

Calì, Massimiliano; Presidente, Giorgio

Working Paper

Robots For Economic Development

Suggested Citation: Calì, Massimiliano; Presidente, Giorgio (2022) : Robots For Economic Development, ZBW - Leibniz Information Centre for Economics, Kiel, Hamburg

This Version is available at:

<https://hdl.handle.net/10419/249581>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Robots For Economic Development

Massimiliano Calì^{*1} and Giorgio Presidente^{†2}

¹*World Bank*

²*Oxford Martin School, University of Oxford*

Abstract

Recent evidence suggests that automation technologies entail a trade-off between productivity gains and employment losses for the economies that adopt them. This paper casts doubts on such trade-off in the context of a developing country. It shows significant productivity and employment gains from automation in Indonesian manufacturing during the years 2008-2015, a period of rapid increase in robot imports. Analysis based on manufacturing plant data provides evidence that the absence of this trade-off is due to diminishing productivity returns to robot adoption. As a result, the benefits from automation could be particularly large for countries at early stages of industrialisation, such as Indonesia. Suggestive evidence indicates these results could apply to developing countries more generally.

Keywords: Robots, Automation, Development, Employment, Productivity.

JEL classification: O14, J23, J24, L11, F63.

* The authors thank Erhan Artuc, Paulo Bastos, Aufa Doarest, Frank Levy, John Morrow, Clemente Pignatti, Bob Rijkers, Rajesh Ramachandran, and Pengfei Zhang for helpful comments, Nabil Ryandiansyah for outstanding research assistance and Mauro Testaverde for sharing automatability data on Southeast Asian countries. The paper also benefited from helpful comments from the OMPTEC discussion group at the Oxford Martin School, and seminars held at Catholic University of Milan, University of Bologna, and Nuffield College, Oxford University. Giorgio Presidente gratefully acknowledges funding from Citi. The views expressed in this paper are those of the authors and do not necessarily reflect those of the affiliating institutions.

† Corresponding author: giorgio.presidente@oxfordmartin.ox.ac.uk

1 Introduction


Ever-more ubiquitous automation technologies present a seeming trade-off for the economies that adopt them. On one hand, they generate significant productivity gains for adopting firms ([Acemoglu et al., 2020](#); [Koch et al., 2019](#)). On the other hand, these gains come at the expense of employment, due to the labor-saving nature of automation. The negative employment impact of automation has been documented for high-income (e.g. [Acemoglu and Restrepo, 2020](#); [Dauth et al., 2019](#)) and a few middle-income economies ([Giuntella and Wang \(2019\)](#) on China and [Artuc et al. \(2019\)](#) on Mexico).

Much less is known about the impact in countries at earlier stages of industrialisation, which include most developing countries today. This is an important gap as the penetration of automation in developing countries is expected to grow significantly over the next decades ([Hallward-Driemeier and Nayyar, 2017](#)). Yet, it is not clear that the existing evidence on automation impact may provide a useful guidance for developing countries. Given that automation technologies may be subject to strong diminishing returns ([Graetz and Michaels, 2018](#)), their adoption may be more likely to increase labor demand in countries at early stages of adoption.¹ Figure 1 provides some suggestive evidence in support of this idea focusing on industrial robots, an important class of automation technology. It shows that the correlation between robot penetration and employment is negative for OECD countries - which have higher robot penetration - and positive for non-OECD countries.

To help fill this gap, this paper examines the impact of robots' adoption on employment in Indonesia, where during the period of analysis the penetration of robots was considerably lower than in the other countries with available micro evidence. The number of robots in the country was very limited before the beginning of our sample in 2008 and accelerated swiftly thereafter. This acceleration was highly heterogeneous across sectors. By the end of the sample period in 2015, the penetration of robots in the most automated industries, such as Motor Vehicles, was similar to that of advanced economies. Other industries, such as Textiles, have seen no

¹ This effect would stand in contrast to the argument of [Diao et al. \(2021\)](#) that the adoption of capital-intensive technology may explain the poor employment performance of manufacturing even in countries at an early stage of industrialization, such as Ethiopia and Tanzania. While their argument is plausible, the data do not allow them to isolate the impact of automation technologies in these countries.

Figure 1: Robot penetration and employment changes: OECD vs non-OECD countries.



The figure plots the correlation between the change in residuals from a regression of log-employment on the share of population above 55 years old over population between 20 and 49 years old, and changes in robot penetration over the same period. Robot penetration is defined as the stock of industrial robots per thousand employed workers. Sources: IFR; PWT; World Bank.

penetration throughout the period. Indonesia has also a rich set of micro- and local-level data on manufacturing plants and employment, enabling us to employ methods comparable to other studies to identify the effects of automation and explore the underlying mechanisms.

In contrast to the available micro evidence in other countries, our analysis documents a positive manufacturing employment effect of robot adoption across Indonesian local labor markets. To identify the effect, we follow the standard approach of constructing a Bartik-style measure (Bartik, 1991) interacting baseline industry shares in a local labor market with industry-specific robot imports (Acemoglu and Restrepo, 2020). We address the endogeneity of robot imports by instrumenting it with a measure of *automation possibilities* based on industry-specific robot penetration in countries ahead of Indonesia in terms of robot adoption (Acemoglu and Restrepo, 2020; Dauth et al., 2019). A large battery of tests provides confidence on the robustness of the positive employment effect of robots.

We then turn to plant-level data to shed light on the micro-mechanisms behind the labor market results. The main analytical challenge is that we do not observe robot adoption at the

plant-level. To overcome this issue, we follow [Autor et al. \(2003\)](#) and assume that industrial robots are best suited to performing routine manual tasks. We document that most workers performing routine task-intensive occupations have secondary school as their educational attainment level, which we observe in the base year. Therefore, we quantify the exposure of plants to robots by interacting yearly imports of robots by 2-digit industry with plants' share of secondary education workers in the base year. We perform a battery of tests to assess the validity of our exposure measure, and show that it is strongly correlated with plant-level investment in machinery and equipment, which includes industrial robots.

Consistent with the predictions of a task-based model (e.g. [Acemoglu and Restrepo, 2018](#)), the micro-level evidence supports the hypothesis of diminishing employment and productivity returns to robot adoption. We find that diminishing returns only affect plants in the top decile of the base-year distribution of robots. As a result, the vast majority of Indonesian plants benefited from particularly large productivity and — hence — employment effects of automation, which explain the positive average effects across the entire distribution.

In the last part of the paper, we consider whether Indonesia is likely a bellwether of future trends in other developing countries, as they move from very low to much higher automation rates. In particular, we examine the possible external validity of our results of positive employment impact and diminishing marginal returns to automation. To that end, we analyze the relation between employment and robot imports across 61 OECD and non-OECD economies in 12 industries over the 2007-15 period. The analysis is based on 2 Stage Least Square estimator instrumenting robot density with a leave-out mean constructed for the same industry-year pairs in other countries. The findings suggest significant diminishing returns to automation across countries. As a result, robot density impacts manufacturing employment negatively in OECD countries — particularly at high level of penetration — and positively in non-OECD countries.

The paper is related to the large literature on firm performance in developing countries (see [Verhoogen, 2020](#) for a review), but it focuses on automation technologies. While an emerging literature studies the impact of automation technologies in advanced and emerging economies ([Acemoglu et al., 2020](#); [Koch et al., 2019](#); [Acemoglu and Restrepo, 2020](#); [Dauth et al., 2019](#); [Giuntella and Wang, 2019](#); [Artuc et al., 2019](#); [Graetz and Michaels, 2018](#); [Humlum, 2019](#);

[Acemoglu et al., 2020](#)), this paper focuses on a developing country. As such, we contribute to the literature on the impact of technology and trade on developing economies, such as [Diao et al. \(2021\)](#) and [Rodrik \(2016\)](#). In particular, our results cast doubts to [Diao et al. \(2021\)](#), who argue that the adoption of capital-intensive technology can explain the poor employment performance of manufacturing, even in countries at an early stage of industrialization, such as Ethiopia and Tanzania.

This paper is also related to the literature examining the impact of automation in advanced economies on developing countries through international trade linkages and reshoring ([Kugler et al., 2020](#); [Artuc et al., 2019](#); [Faber, 2018](#); [Artuc et al., 2018](#)). In contrast to that literature, this paper provides novel evidence on the impact of automation in a developing country *on its own economy*.

The rest of the paper is organized as follows: Section 2 presents the positive long-run impact of robots on local labor markets in Indonesia. Section 3 examines candidate mechanisms underlying the aggregate impact, decreasing returns to automation; Section 4 explores decreasing returns to automation using plant-level data; Section 5 frames the results in a broader context and speculates on their external validity, and Section 6 concludes.

2 The Positive Employment Impact of Robots on Local Labor Markets

This section provides evidence of the labor market impact of industrial robots in a developing country — Indonesia — during a period of rapid increase in robot imports.

2.1 Robots and Local Labor Market Data

To do so, we match imports of industrial robots from the International Federation of Robotics (IFR) and Indonesian labor force survey data on employment in local labor markets (Survei Tenaga Kerja Nasional—Sakernas).

Industrial robots are defined by the International Standards Organization as *an automatically controlled, reprogrammable, multipurpose manipulator programmable in three or more*

axes, which can be either fixed in place or mobile for use in industrial automation applications.² The autonomy and physical dexterity typical of robots make them particularly suitable to perform tasks otherwise performed by human workers. It is precisely their labor-saving nature that sets industrial robots apart from other categories of capital assets.

The IFR collects data from each national robotics association. Since almost all robot suppliers are members of national associations, the dataset includes virtually all robots used worldwide. An advantage of the data is that the IFR has a common protocol to count robots, so that it ensures consistency across countries and years. Information is available for each country, 2-digit industry and year. Details about the industry-level data on robots are presented in Online Appendix B.1.

The data shows that the total number of robots shipped to the Indonesian manufacturing sector was low and roughly constant up to 2009, when the country had a stock of approximately 350 units in use (Figure 2).³ Since 2010, the pace of increase in robots import experienced a significant uptick with the total stock increasing almost 10-fold by 2015. Our analysis exploits this large jump in adoption of robots to identify their labor impact.

The building block of the empirical methodology is industry-level robot penetration, defined as the number of industrial robots shipped to industry i (measured at 2-digit ISIC code) in year t , divided by the number of workers in industry i (in thousands):

$$\frac{R_{i,t}}{L_{i,t_0}}$$

To relieve potential endogeneity concerns, we fix the number of industry workers (the denominator) to the base year value, in the same vein as [Acemoglu and Restrepo \(2020\)](#).


The aggregate trend presented in Figure 2 hides a large amount of industry heterogeneity in the use of robots. Figure 3 shows that the number of robots per thousand workers used in manufacturing is much higher in Motor vehicles and Rubber and plastics than in other industries.⁴

² ISO 8373:2012.

³ Indonesia does not produce robots domestically. Therefore, IFR data coincide with total imports of robots from Comtrade. However, unlike IFR data, Comtrade does not provide an industry breakdown.


⁴ While Motor vehicles is by far the most automated industry worldwide, the high concentration of robots in Rubber and plastics is more peculiar to Indonesia. This is one of the largest manufacturing industries in Indonesia particularly for auto-parts, such as tires, which often employ state-of-the-art technologies of production.

Figure 2: Total number of robots used in the Indonesian manufacturing sector.



The figure shows the total number of industrial robots used in the manufacturing sector over the years of the sample. Source: IFR.

Figure 3: Penetration of robots in the Indonesian manufacturing sector, by industry.



The figure shows the number of industrial robots per thousand workers used in selected industries over the years of the sample. Source: IFR, SI.

Labor force survey data comes from the Indonesia National Statistics Bureau (BPS), which collects data on a large cross-section of workers bi-annually.⁵ Regencies represent the second level of sub-national administrative divisions in Indonesia, the first being provinces. A number of features make them a reasonable proxy for local labor markets in Indonesia. First, the mobility of labor is limited across regencies. In 2010 for example, Sakernas data shows that only 5 percent of the workforce worked in a different regency than their residence.⁶ Second, following Indonesia’s 1999 decentralization reforms, regencies hold significant administrative powers, including in the labor markets, such as the minimum wage setting. We use the pre-decentralization reform division in 292 regencies to ensure the consistency of the analysis over time.

The details of the labor market data are provided in Section 2.1. Appendix Table I presents the summary statistics.

2.2 Local Labor Market Analysis: Empirical Approach

To identify the effects of robots on local employment, we follow the standard approach of constructing a Bartik-style measure (Bartik, 1991) interacting baseline industry shares in a local labor market with industry-specific robot imports (Acemoglu and Restrepo, 2020; Dauth et al., 2019; Artuc et al., 2019; Giuntella and Wang, 2019).

We set 2007 as the base year as Sakernas data is representative at the regency-level only since 2007. For each industry in the base year, we compute industry i share of employment in regency r total employment:

$$s_{i,r} = \frac{L_{r,i,2007}}{L_{r,2007}} \quad (1)$$

we then use these shares to aggregate industry-specific changes in robot penetration weighted by the initial regency’s industry composition. Accordingly, we define Indonesian regency-level

⁵ We focus on wage employees as those are the workers more directly affected by firms’ adoption of robots. This focus is also relevant from a policy perspective, as wage employment typically provides a more reliable and higher incomes than self-employment in a developing country like Indonesia.

⁶ Indonesia’s island geography and often underdeveloped transportation infrastructures make the hypothesis of limited mobility likely to hold.

exposure to changes in robot penetration as follows:

$$\Delta ETR_r^{ID} \equiv \sum_{i \in r} s_{i,r} \frac{R_{i,2015}^{ID} - R_{i,2008}^{ID}}{L_{i,2007}^{ID}} \quad (2)$$

An obvious problem with using ΔETR_r^{ID} is its endogeneity to local employment changes, as unobserved industry-level shocks may be correlated with both robot adoption and changes in labor market outcomes. For instance, if investment in robots is driven by the need of keeping up with demand in response to a positive shock, the estimated impact of ΔETR_r^{ID} on employment would be upwardly biased.

As in other studies, we address this issue by developing a measure of *automation possibilities* to instrument *actual automation adoption* (Acemoglu and Restrepo, 2020; Dauth et al., 2019). We focus on industry-specific robot penetration in countries that are ahead of Indonesia in terms of robot adoption. OECD countries appear the most suitable reference group for a middle-income country as Indonesia. We match IFR data with 2-digit industry employment figures from the OECD Structural Analysis Database (STAN). For each industry-year pair, we compute the number of imported robots per thousand workers, averaged across OECD countries. Then, the measure of automation possibilities is computed as in (2), but replacing $\frac{\Delta R_{i,t}^{ID}}{L_{i,t}^{ID}}$ with the OECD counterpart.


If robots adoption were driven by exogenous improvements in technology, e.g. a reduction in robot prices, then we should observe a positive correlation between adoption in Indonesia and in other countries.⁷ Figure 4 confirms a strong positive correlation between 2008-2015 changes in robot penetration in Indonesia and in OECD countries.⁸ This pattern is consistent with the hypothesis that Indonesia's automation across industries is driven by global shocks, rather than specific domestic conditions. Our identification is based on the assumption that such shocks to automation possibilities affect domestic trends only through their impact on actual

⁷ For instance, an enhanced machine-vision technology applicable to painting auto parts, should spur robot adoption in the motor vehicles industry worldwide.

⁸ As the change for Motor vehicles and Plastic and rubbers is much larger than for the other industries, the figure uses a log-scale to ease readability. The log-scale allows to include only nine of the twelve available industries as Textile and Paper did not experience any robot adoption in Indonesia and the average change in robots per thousand workers was negative in Other transport equipment in the OECD region. Online Appendix Figure A1 shows that the positive relationship holds also when using a normal scale, which allows to include Textile, Paper and Other transport equipment, while it excludes Motor vehicles and Plastic and rubbers.

robot use in Indonesia. In Section 2.4 we test the robustness of this assumption to a variety of checks, including testing for the exogeneity of the shift-share instrument and the potential impact of non-robot technologies.

Figure 4: Correlation between Indonesian and OECD-average exposure to robots (log-scale).



On the horizontal axis there is the change between 2008 and 2015, of the OECD region industry-average number of robots per thousand workers. On the vertical axis, there is the change between 2008 and 2015, of the industry-level number of robots per thousand workers in Indonesia. Sources: IFR, STAN, SI.

Our baseline results are derived from 2SLS estimates using the following system of equations:⁹

$$\begin{cases} \Delta ETR_r^{ID} = b_0 + b_1 \Delta ETR_r^O + \alpha \Delta D_r + u_p + \epsilon_r \\ \Delta E_r = \beta_0 + \beta_1 \widehat{\Delta ETR_r^{ID}} + \gamma \Delta D_r + u_p + \epsilon_r \end{cases} \quad (3)$$

where ΔETR_r^O is the instrument based on OECD penetration and ΔE_r is the change in employment between 2008 and 2015 in regency r (in logs).¹⁰

In line with the rest of the literature on automation and local labor markets, ΔETR_r^{ID} takes

⁹ We also show that our results are very similar if we estimate OLS reduced forms.

¹⁰ All employment changes in our dataset are positive and so we do not face issues of zero or negative values.

the form of a shift-share variable.¹¹ The identification of the parameter of interest, β_1 , arises from variation in two components: the “shift” in industry robot penetration and the “share”, i.e. differences in local exposure to industries hosting robots.

Our identifying strategy relies on the assumption that no other unobservable factors are correlated with both automation opportunities and changes in employment at the regency level. The first difference transformation of equation 3 absorbs all such possible factors whose effect is time-invariant, e.g. geography and factor endowment. However, provinces have labor market powers in Indonesia, including minimum wage and the enforcement of labor regulation. Differences in regulation might induce differential trends, in turn affecting initial industry composition and employment changes. For this reason, we include province fixed effects, u_p , in (3).

An important concern is that *global* industry shocks could result in a violation of the exclusion restriction. For instance, a positive global demand shock to an industry could increase employment in an Indonesian regency hosting that industry. It could also enhance robot adoption in that industry (including in the OECD), aimed at keeping up with the higher demand. To mitigate such concerns, we construct a regency-level demand shifter, ΔD_r and include it in (3). The shifter aggregates global exports by industry (excluding Indonesian industries) at the regency-level, using the shares in (1). In Section 2.4, we also test the robustness of the results to controlling for a variety of other baseline covariates, which might be related to time-varying shocks correlated to both automation possibilities and changes in employment.

We cluster standard errors at the regency-level. We also experiment below with alternative clustering and calculate shift-share standard errors, following the methodology of [Adao et al. \(2019\)](#).¹² We weigh the estimate by regencies’ population in the base year so as to obtain nationally representative estimates.

¹¹ See Section 2.4 for a discussion of the conditions for consistent identification with a shift-share variable in our context and a battery of diagnostic tests.

¹² [Adao et al. \(2019\)](#) show that shift-share designs, such as those used in this part of the analysis, might lead to residuals that are correlated across regions with similar industry shares and so over-rejection of the null hypothesis of no impact. In our framework, this is important because if there are omitted variables correlated to industry composition and employment changes, these would enter the error term and introduce correlation across regencies.

2.3 Local Labor Market Results

Column (1) of Table 1 presents the first stage estimates, which confirm the strength of the instrument. Regency-level exposure based on OECD robot penetration is a strong predictor of Indonesian exposure. The coefficient is significant at the 1 percent level and the regression explaining over half of the overall variation in exposure.

The dependent variables in columns (2)-(5) are regencies' employment changes. Column (2) presents the estimated β_1 coefficient for total employment, which is positive, although not statistically significant. We then focus on manufacturing, which is the robot-adopting sector. The estimates in column (3) suggest that robots have a positive long-run impact on manufacturing employment, significant at the 1 percent level. The magnitude of the coefficient implies that employment grew 31 percentage points more in regencies with one additional robot per thousand base-year workers over the 2008–15 period.¹³

This result is in contrast with findings based on advanced and emerging economies. These suggest that industrial robots have a negative impact on local level employment — most notably in manufacturing, which is the sector employing robots in production.¹⁴ Much of the remainder of the paper will focus on testing and explaining the robustness of this important discrepancy.

Columns (4)-(5) present the results of the impact of robots on other non-manufacturing sectors. The coefficient for services is positive, but small and not statistically significant (column 4). Conversely, the impact of automation on employment in primary sectors (agriculture and mining), is negative and statistically significant (column 5). The results suggest a reallocation of labor towards manufacturing as a result of automation in manufacturing. This accelerates the process of workers moving out of agriculture in particular, as mining represents a small share of workers in primary sectors in Indonesia.

¹³ Alternatively, a doubling in the number of robots relatively to the mean (0.16) would generate an increase in manufacturing employment growth of 5 percentage points.

¹⁴ [Acemoglu and Restrepo \(2020\)](#) and [Dauth et al. \(2019\)](#) find evidence of a negative impact on employment in the United States and Germany, respectively. [Giuntella and Wang \(2019\)](#) and [Artuc et al. \(2019\)](#) provide similar evidence for China and Mexico, respectively. The literature also points to a negative impact of robots on average wages in manufacturing. Results in [Calì and Presidente \(2021\)](#) suggest that exposure to robots has instead a positive impact on average wages in Indonesian manufacturing.

Table 1: Regency-level results: employment in manufacturing and non-manufacturing sectors.

	(1) Regency ETR	(2) Δ Total employment	(3) Δ Manufacturing employment	(4) Δ Services employment	(5) Δ Agr/Mining employment
Regency ETR (instrument)	0.440*** (0.082)				
Regency ETR		0.054 (0.039)	0.309*** (0.105)	0.023 (0.057)	-0.277** (0.137)
Observations	276	276	276	276	276
R-squared	0.472				
Province FE	yes	yes	yes	yes	yes
Regency demand shifter	yes	yes	yes	yes	yes
First stage F-stat	—	28.97	28.97	28.97	28.97

The table presents 2SLS estimates of the relationship between regency-level exposure to robots and employment. Indonesian exposure to robots is instrumented with exposure based on OECD industry penetration. The dependent variables are the 2008-2015 differences of log of employment in each regency. The regency demand shifter aggregates global exports by industry (excluding Indonesia) using regency-level industry employment shares in the base year. Standard errors are clustered at the regency-level. Weights are constructed using 2007 (base year) regency population. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

2.4 Threats to the Validity of the Estimates

We start by checking the robustness of one of the model’s key assumptions, i.e. that initial sectoral specialization is not correlated to unobserved regency-level factors with an impact on employment changes. Online Appendix Figure A2 plots the raw correlation between regency-level manufacturing employment changes and base-year regency industry shares. Reassuringly, the two variables appear to be largely uncorrelated in our sample.

Given that the largest changes in robots penetration are concentrated in few industries (figures 3 and 4), the identification of β_1 should arise mainly from local differences in exposure to a few common shocks.¹⁵ We follow the approach in Goldsmith-Pinkham et al. (2020) and compute the so-called “Rotemberg weights”, i.e. relative importance of each industry in determining the overall explanatory power of the shift-share instrument. As in Goldsmith-Pinkham et al. (2020), we find that a small number of industries have a large share of the weight. In our sample Food products and Basic metals together account for over 95 percent of the positive weight of the estimator.¹⁶

¹⁵ Borusyak et al. (2018) for an alternative interpretation based on the assumption of quasi-random variation in shocks across many industries.

¹⁶ In their application, Goldsmith-Pinkham et al. (2020) find that only 5 of the available 228 industries explain almost half of the predictive power of the Bartik variable. Therefore, it is not surprising that in our sample with 13 industries, only two explain most of the weight.

We then probe the validity of our identification assumption focusing on such two industries. Online Appendix Table A1 regresses the regency-level shares of the Food products industry (column (1)) and Basic metals (column (2)) on several base-year covariates that might affect industrial composition as well as changes in employment. In column (3), we repeat the exercise for the ETR instrument as a whole. We do find that some covariates enters the regression significantly and so we re-run the baseline specification including such controls. Reassuringly, however, Online Appendix Table A2 shows that the results are virtually unchanged. Our results are unchanged if we exclude Food products and Basic metals, or the motor vehicles industry, which experienced the largest increase in robot penetration over the sample.¹⁷

We test further the validity of the exclusion restriction by providing evidence that automation possibilities are uncorrelated to other trends affecting Indonesian industries. In particular, we are concerned that the adoption of computers or other non-automation technologies, potentially complementary to robot use, might be driving the positive employment impact. To that end, we estimate industry-level correlations between OECD robot penetration and an index of technological sophistication including information on IT usage, R&D, product and process innovation.¹⁸ The results are presented in Online Appendix Figure A3. In Panel (a), we correlate the change in industry-level OECD penetration between 2000 and 2006 (in logs), to the industry average value of the technological index in 2006.¹⁹ The figure shows that there is no correlation between the two variables. Panel (b) presents the same relationship for capital intensity, measured as the average capital-labor ratio (in logs). In this case, we find a mild but positive relationship between the two variables. Panel (c) shows instead the correlation between 2008-2015 changes of OECD robot penetration and changes in the capital stock over the same period. Also in this case, the correlation is essentially flat.

In order to further mitigate the concern that our instrument might be correlated to the adoption of non-robot technologies, we check whether the 2008-2015 changes of OECD robot pene-

¹⁷ In fact, if we exclude Food products and Basic metals, we find a positive impact of robots on total employment as well. The tables are available upon request.

¹⁸ The index is based on plant-level data. See Section 4.3 for details on its construction.

¹⁹ Ideally, we would have wanted to correlate changes in OECD penetration to changes in computer usage. However, Sakernas labor market data do not provide information on technology usage. Therefore, we exploit plant-level data to calculate industry averages. Since information on technology usage is only available for the census year 2006, the most sensible way of presenting the evidence is correlating *past* changes in OECD robot penetration to 2006 values of the index.

tration is correlated with changes in capital intensity across regencies. The results are presented in column (1) of Online Appendix Table A3. Reassuringly, we find no evidence of correlation between the two variables. In column (2), we regress the log-change in the aggregate capital stock on the instrument. The coefficient is positive, but still not statistically significant.

Next, we verify that changes in manufacturing employment between 2000 and 2007 do not predict our instrument, nor the regency-level shares of industries with the largest Rotemberg weights, Food products and Basic metals.²⁰ This is shown in columns (1)-(3) of Online Appendix Table A4. None of the coefficients is statistically significant, which reassure us about the absence of significant pre-trends. Before turning to the robustness of the labor market estimates of Table 1, we show graphically that regency-level employment changes between 2000 and 2007 were not correlated to employment changes between 2008 and 2015. The results are presented in Online Appendix Figure A4, which shows no sign of correlation between employment growth across the two period.

The labor market results in Table 1 are robust to several alternative specifications. First, we check if our main result depends on whether we estimate the model with OLS or 2SLS. Online Appendix Table A5 presents reduced-form estimates using the OECD-based exposure instrument in an OLS regression. The coefficients are smaller than the 2SLS estimates, but the signs and significance of the coefficients are analogous to those of the baseline.

Second, we estimate a model in stacked-differences by splitting our sample between the years 2008-2009 and 2010-2015, which Figure 3 shows to exhibit different trends in robot penetration. Online Appendix Table A6 shows that the magnitude of the coefficients is largely unaffected, but they are more precisely estimated than the baseline in long-differences.

Third, automation might be not the only technological force affecting Indonesian employment. In particular, other forms of routine-biased technical change (Autor et al., 2003; David and Dorn, 2013) might be driving our estimates. To test whether that is the case, we rely on data from Frey and Osborne (2017), which provide a list of occupations in manufacturing and non-manufacturing industries at high risk of automation.²¹ We use the list to calculate the share

²⁰ Since Sakernas labor market data are not fully representative before 2007, we pool bi-annual 2000 and 2001 waves to maximise representativeness.

²¹ Specifically, Frey and Osborne (2017) calculate the probability of computerisation based on the description of the tasks performed in each occupation. Then, they deem an occupation at risk of automation if the associated

of Indonesian employment at risk of automation in each regency in the base year.²²

To proxy for the impact of technology more broadly, we control for the base-year value of the regency-level average of the technological sophistication index, and the log of the capital-labor ratio. The results obtained adding the three variables are presented in Online Appendix Table A7. The effects of automation on aggregate and sectoral employment remain very similar to the baseline specification. Interestingly, we find that a large share of routine employment is associated to negative employment changes in manufacturing, but positive changes in services. This is consistent with a reallocation of routine workers from manufacturing to low-skill services (David and Dorn, 2013). More technologically sophisticated regencies tend to have larger changes in total employment, but this is driven by employment gains in services. We do not find any significant impact of regency-level capital intensity on employment.²³

Finally, in Online Appendix Table A8, we experiment with alternative computation of standard errors. In round brackets, we present standard errors clustered at the province-, rather than regency-level. This allows for a correlation of the errors at another administrative level, which holds some policy powers relevant to labor markets. Province-level clustered errors tend to be slightly larger, but the coefficients are still statistically significant. Squared brackets include shift-share standard errors, obtained with the methodology in Adao et al. (2019).²⁴ The coefficients on manufacturing employment and agriculture are still statistically significant with shift-share errors.

probability is greater than 0.7. The calculation of the probabilities relies on information from the 2010 version of the Dictionary of Occupational Titles (ONET), an online tool developed for the US Department of Labor.

²² Besides information on employment status, sector and work location, Sakernas also includes information on the occupation of the worker identified according to the Indonesian classification (Klasifikasi Baku Jenis Pekerjaan Indonesia—KBJI). This is compatible with the International Standard Classification of Occupations (ISCO), which can be mapped to the data from Frey and Osborne (2017).

²³ This is not surprising, as Online Appendix Figure A3 shows that such variables are at best mildly correlated with robot penetration at the industry-level.

²⁴ Specifically, we follow Dauth et al. (2019) and impose the null hypothesis $H_0 : \beta_1 = 0$. Adao et al. (2019) show that the applying such adjustment improves the finite-sample properties of the confidence intervals when the number of industries is relatively small. The computation is performed with the user-written STATA command *reg_ss* (Adão et al., 2020).

3 Diminishing Returns from Automation


As discussed above, to the best of our knowledge this is the first empirical analysis finding a positive local labor market impact of robots on employment. What can explain the difference with the empirical results from other countries?

To answer this question, it may be useful to recall some suggestive evidence of diminishing productivity returns to robots unveiled in [Graetz and Michaels \(2018\)](#). That evidence is based on a sample of high-income countries in 1993-2017, when their robot density was considerably higher than in Indonesia during our period of analysis.

In fact, robot penetration in Indonesian manufacturing was much lower than in all other economies for which labor market evidence on robots exists. Figure 5 illustrates the extent of the discrepancy. For each of these countries, the figure plots the number of robots per million manufacturing workers in the first (if available to us) and last years of the respective analysis. The data for robots comes from IFR and for manufacturing employment from the ILO and — for China — from UNIDO. For the missing years, we input the manufacturing employment values on the basis of the share of industry in employment and of total employment in those years (again from the ILO). Take the last years of the analyses first, as these are close in range for most countries (2014-2016 except the US, which refers to 2007). Indonesia exhibits the lowest robot penetration among the sample of countries by a factor ranging between 9 (vis-à-vis Mexico in 2015) and 99 (vis-à-vis Germany in 2014). Similar differences apply also when taking the first years of the analyses. Indonesian robot density relatively to the other countries varies between a factor of 4 (vis-à-vis China in 2006) and 158 times smaller (vis-à-vis Germany in 1994).

To the extent that diminishing productivity returns to automation are at play, this could translate in increases in labor demand at very low densities of robots, as in the case of Indonesia. This could also explain the much less negative employment impact of automation in Mexico ([Artuc et al., 2019](#)) relative to the US ([Acemoglu and Restrepo, 2020](#)). We illustrate this idea using a task-based model along the lines of [Acemoglu and Restrepo \(2018\)](#).

Figure 5: Robot penetration in Indonesian vs other countries with evidence on the impact of automation.



The figure shows the number of robots per million manufacturing workers in the first (if available to us) and last years of the respective countries' analysis. These are: 1994-2014 (DEU); 1990-2015 (ESP); 2010-15 (FRA); 1993-2007 (USA); 2006-16 (CHN); 2000-15 (MEX); 2008-15 (IDN).

Sources: IFR; ILO; UNIDO.

3.1 A Simple Task-Based Model

Consider an economy in which the unique final good is produced by a representative firm combining a continuum of varieties with total mass equal to 1. The parameter $\sigma > 1$ governs the constant elasticity of substitution across varieties.

Each differentiated variety is produced by a single intermediate good-producing firm f , which combines a unit measure of tasks. Let $x_f(z)$ represents the production services of task z . The output of firm f is given by

$$y_f = \exp \left(\int_0^1 \ln x_f(z) dz \right)$$

Since all firms are identical and face the same problem, we suppress f subscripts and focus on a representative firm.

Task services are produced using machines and human labor, which are perfectly substi-

tutable. Markets for both inputs are perfectly competitive. However, humans are more productive at performing certain tasks. Let $-\infty < \tilde{\gamma}(z) < \infty$ be the log-productivity of labor relative to machines in performing task z , which is the same across firms.²⁵ Without loss of generality, we assume that $\tilde{\gamma}(z)$ is increasing in z , so that higher-ordered tasks are harder to automate.²⁶ Let κ denote the given extent of firm-level automation, such that tasks $z \in [0, \kappa]$ are performed by machines and tasks $z \in (\kappa, 1]$ by labor. In partial equilibrium, firms take factor prices are given. Let n represent labor demand. Online Appendix D shows that the employment effect of automation is increasing in the productivity effect $\left(\frac{\partial \ln y}{\partial \kappa}\right)$ and decreasing in the displacement effect $\left(\frac{1}{1-\kappa}\right)$:

$$\frac{\partial \ln n}{\partial \kappa} = \frac{\sigma - 1}{\sigma} \frac{\partial \ln y}{\partial \kappa} - \frac{1}{1 - \kappa} \quad (4)$$

Let ω be the log price of labor relative to capital. Online Appendix D shows that:

$$\frac{\partial \ln y}{\partial \kappa} = \sigma(\omega - \tilde{\gamma}(\kappa))$$

The productivity effect of automation for firm f is increasing in ω , because substituting labor with machines reduces marginal costs more when labor is relatively expensive. However, the productivity effect is lower for firms in contexts where a larger number of tasks is already automated (i.e. κ is larger):

$$\frac{\partial^2 \ln y}{\partial \kappa \partial \kappa} = -\sigma \tilde{\gamma}'(\kappa) < 0$$

This arises from the fact that the $\tilde{\gamma}(\cdot)$ is increasing, which reflects the comparative advantage of humans in performing non-routine, complex tasks. In a context like Indonesia during our period of analysis, where only a relatively small share of tasks is automated, there are more opportunities for productivity gains. On the contrary, in a context where producers have exploited most available opportunities to automate, using machines to perform the marginal task

²⁵This assumption is more likely to hold for firms operating in industries characterised by similar production processes. To account for this explicitly, one possibility is assuming that the final good is industry-specific and final consumers maximise utility over a continuum of final goods. However, this would complicate the model without altering our conclusions.

²⁶ For instance, these could be non-routine tasks requiring creativity (Autor et al., 2003).

would deliver limited gains. That is because of the strong comparative advantage of humans in performing complex tasks that are not yet automated.²⁷

The displacement effect, on the other hand, is always negative and monotonically decreasing in κ . This is due to the fact that as the range of tasks performed by humans shrinks due to automation, the marginal product of labor decreases and workers become redundant. The previous conditions imply that the employment effect is also decreasing in the initial level of automation:

$$\frac{\partial^2 \ln n}{\partial \kappa \partial \kappa} = -(\sigma - 1)\tilde{\gamma}'(\kappa) - \frac{1}{(1 - \kappa)^2} < 0$$

Thus, for countries where automation is already widespread, increased automation is more likely to decrease employment not only because the productivity effect is lower, but also because the displacement effect is stronger. This happens because if a wide range of tasks are already automated, the labor displaced by the additional automation has a low marginal productivity given the marginal decreasing returns to labor with respect to the range of activities.

The task-based model thus suggests that the positive employment impact of automation presented in Section 2 might reflect a large productivity effect offsetting the labor displacement effect. This is more likely to happen in a country where firms tend to be less automatized, as in Indonesia.

4 Diminishing Returns in Plant-level Data

This section tests these relationships with Indonesian plant-level data. It describes in turn the data, methods and results of the empirical analysis.

4.1 Plant-level Data

The main data source for the plant-level analysis is the Indonesian survey of manufacturing plants with at least 20 employees (Statistik Industri, SI henceforth). This is an extensive panel

²⁷ Acemoglu and Restrepo (2019) refers to technologies conferring only limited productivity gains as “so-so technologies”.

of plants, which becomes a census at ten-year intervals and covers over 90 percent of the relevant population in the inter-census years. It includes detailed data on production and sales, including quantities produced by products, on employees, costs and assets, and specifically in machinery and equipment (which includes industrial robots). Online Appendix B.2 provides more detailed information about the plant-level data .

We match SI and IFR data by constructing industry crosswalks, which are presented in Online Appendix B.3.²⁸ Around 8.5 percent of the plants in our sample switch industry over the period of analysis. To avoid potential endogeneity of industry choice to robot adoption, we drop the observations after a plant switches industry. We obtain very similar results if we do not drop such observations, or adopt alternative strategies.²⁹

Our final plant level (unbalanced) panel includes 15,351 plants operating in 13 2-digit manufacturing industries between 2008 and 2015, for a total of 65,573 observations.³⁰ Appendix Table II presents summary statistics for the variables used in the plant-level analysis.

4.2 Measuring Plant-Level Exposure To Robots

Unlike some other studies in high-income countries (e.g. [Acemoglu et al., 2020](#) and [Koch et al., 2019](#)), we do not observe the use of robots by plants. Instead, we match data on Indonesian imports of industrial robots by industry with plants' observable characteristics to build a measure of plant-level exposure to robots.

The key assumption underlying our proxy of plant exposure is that routine task-intensive occupations are the most likely to be automated ([Autor et al., 2003](#)). As we do not observe the tasks performed by workers in each plant, our baseline measure of exposure proxies routine-intensive labor with the educational level of each plants' workforce, which we observe in 2006 when the manufacturing plants' survey becomes a census.

²⁸ Industry classifications in both SI and IFR data are roughly equivalent to ISIC Rev. 4. However, in some cases SI industries are more granular than IFR. Thus, we group together some SI industries to ensure maximum compatibility across the two datasets.

²⁹ The first alternative strategy is assigning robots to plants based on the industry in which it operates the year of the first available information. The second strategy is dropping altogether all switching plants. Results of these alternative estimations are available from the authors upon request.

³⁰ In principle, we could run the plant level analysis starting in 2007. However, representative data by regency and industry, which are discussed in Section 2.2, are only available from 2008. Therefore, to be consistent we start the plant-level analysis in 2008. Plant level results for the period 2007-2015 are very similar, which is not surprising given that robot adoption started to take off in 2009 (see Figure 2).

Specifically, we identify the production occupations most exposed to being automated using the definition of occupations’ “replaceability” from [Graetz and Michaels \(2018\)](#). Then, using Indonesian labor force survey data (described in Section 2.1), we compute the distribution of workers in occupations at high risk of automation by their educational attainments in 2006, the first year such information is available. Table 2 presents the result and shows that workers in occupations at risk of automation in Indonesia are typically secondary education workers.³¹

Table 2: Share of employed workers in occupations at high risk of automation, by educational attainment.

Educational attainments:	Primary	Secondary	Tertiary
Share of employment at high risk of automation	.39	.59	.02

The table reports the share of employment in production occupations at high risk of automation, by the educational attainments of Indonesian workers in 2006. Occupations at risk of automation are identified using the methodology of [Graetz and Michaels \(2018\)](#). Primary education includes up to completed primary school. Secondary education includes junior and senior high-school. Tertiary education includes education levels from diplomas to PhD. Sources: Sakernas; [Graetz and Michaels \(2018\)](#).

The finding that occupations at risk of automation are dominated by secondary educated workers is consistent with various pieces of evidence. First, the literature on employment polarization suggests that automation technologies tend to replace occupations with an intermediate level of skills.³² While we do not observe the skill level of plants’ workforce directly, secondary education is both the median and the mean level of education of workers in manufacturing plants in 2006.

Second, an analysis of the characteristics of production occupations compiled by the World Bank for Indonesia, Thailand and Malaysia, suggests that secondary education is the typical educational attainment for all of the occupations at greater risk of automation across countries with similar economic characteristics (see Online Appendix C for details on the Occupation Profiles).³³

Third, industries with an initial large share of secondary education workers have adopted

³¹ Similar results are obtained using the probabilities of computerization calculated by [Frey and Osborne \(2017\)](#) (see [Cali and Presidente, 2021](#)). However, while [Frey and Osborne \(2017\)](#) consider production and non-production occupations, the notion of replaceability in [Graetz and Michaels \(2018\)](#) is more closely related to industrial automation. Given the focus of the paper on robots, using the list of occupations from [Graetz and Michaels \(2018\)](#) seems more appropriate in our framework.

³² E.g. [Goos et al. \(2014\)](#); [David and Dorn \(2013\)](#); [Goos et al. \(2009\)](#); [Goos and Manning \(2007\)](#).

³³ We thank Mauro Testaverde for sharing these unpublished profiles with us.

relatively more robots in subsequent years (Online Appendix Figure A5). In addition, Online Appendix Figure A6 shows that industry averages of investment in machinery and equipment from SI, which includes investment in robots, is positively correlated to industry imports of robots from IFR data.

Formally, we define exposure to robots (ETR) for plant f as:

$$ETR_{f,t} \equiv \frac{R_{i,t}}{L_{i,t_0}} \times secondary_{f,t_0} \quad (5)$$

where $secondary_{f,t_0}$ is the plant-level share of workers with secondary educational attainments in the base year. Plants that at the beginning of the sample had a large share of secondary education workers are expected to be more exposed to robots as industry penetration grow. On the contrary, ETR is zero if either a plant operates in an industry that does not import robots, or if $secondary_{f,t_0} = 0$ and so the plant does not have opportunities for automation.

To probe further the validity of our assumption that secondary education proxies for routine task-intensity in Indonesian manufacturing, we regress plant-level investment in machinery and equipment in every year on plant-specific ETR controlling for various plant-, year- and industry-year specific controls.³⁴ The result is presented column (1) of Appendix Table III. Reassuringly, the ETR coefficient is positive and statistically significant. It suggests that one additional robot per thousand workers in a given industry is associated with a nine percent increase in machinery and equipment investment by plants one standard deviation above the average share of secondary education workers in that industry. In columns (2) and (3), we report the results of two placebo tests. The first shows that investment in machinery and equipment is *inversely* related to the share of primary education workers — typically performing non-routine manual tasks that are hard to automate. The second shows that these investments are not correlated with the share of workers with tertiary education, consistently with the idea that they tend to perform complex non-routine cognitive tasks (Autor et al., 2003).

³⁴ Specifically we employ (6), which is described in the next section, replacing the dependent variable with investment in machinery. Given the lumpiness of investment at the plant level, for the dependent variable we employ the log-transformation $\ln[x + (x^2 + 1)^{.5}]$. We obtain qualitatively identical results using a dummy variable taking value 1 if, in a given year, a plant has positive investment.

4.3 Plant-Level Analysis: Econometric Specification

We use *ETR* to identify the impact of automation on plant-level outcomes by estimating the following linear model:³⁵

$$Y_{f,i,t} = \gamma_0 + \gamma_1 ETR_{f,t} + \Gamma X_{f,t} + \eta_f + u_{i,t} + \epsilon_{f,t} \quad (6)$$

where $Y_{f,i,t}$ is an outcome of plant f in industry i at time t ; $X_{f,t}$ is a vector of controls, to be discussed further below; η_f are plant fixed effects, and $u_{i,t}$ are industry-year effects defined using the same classification as IFR robot import data.

Plant fixed effects absorb the confounding effect of any plant-level unobserved time-invariant characteristic. This includes factors that might be related to plants having a large base-year share of secondary education and their outcomes. We analyse this aspect in detail further below.

Industry-year fixed effects mitigate the concern that exposure to robots in an industry may be endogenous to the outcomes of plants operating in that industry. They capture any industry-specific, time varying shock that may be related to both the outcome variable and robot penetration, such as changes in international trade patterns, or demand and supply shocks. At the same time, $u_{i,t}$ absorbs the variation of industry-specific robot penetration, $\frac{R_{i,t}}{L_{i,t_0}}$, which therefore cannot be estimated.

Estimates of γ_1 in (6) quantify the *net* impact of one additional robot in exposed plants (i.e. high levels of $secondary_{f,t_0}$), relatively to less exposed ones (i.e. low levels of $secondary_{f,t_0}$) within the same industry. A possible threat to the validity of the estimate is externalities from treated to control plants (Miguel and Kremer, 2004). For instance, an increase in robot penetration in industry X raising productivity of exposed plants in X could also benefit less exposed plants in the same industry through an increase in demand for inputs. The industry classification used by IFR data is sufficiently broad that most inputs are typically sourced within the

³⁵ The plant-level approach is different from the local labor market analysis in that it exploits yearly variation in robot penetration. We argue that the yearly analysis is appropriate for two reasons. First, running the plant level specifications in long difference would reduce the sample size by almost two-thirds relatively to the yearly regression model, as it would only focus on surviving plants throughout the entire period. Labor force survey data are not subject to this problem, which allows us to estimate models in long differences. Second, dropping plants that were in the sample in 2006 but exited before 2015, might introduce selection bias.

same industry. Hence this increase in demand could translate into higher sales and employment for upstream producers within the same industry regardless if they are automated.³⁶ This would result in downward-biased estimates of γ_1 . In order to obtain a clean control group, we exploit 9-digit-level information on inputs purchased by each plant and obtain a detailed input-output table (see Online Appendix B.5). We then use this table to compute downstream exposure to robots in each 5-digit industry and include it in $X_{f,t}$.³⁷ The details of the construction of the downstream exposure variables are presented in Online Appendix E.

Similarly to Section 2.2, we are also concerned that the adoption of other non-automation technologies, such as computers or software, might bias our estimates. The bias would be particularly severe if adoption were correlated to the share of secondary education workers. A related concern is that skill-biased technical change might exhibit complementarity with routine labor, which would result in unobserved trends in technology adoption that might bias the estimates. To mitigate such concerns, we exploit plant-level information available in 2006 (census year) on R&D units, product and process innovation, use of computers and the Internet to construct a plant-level index of technological sophistication.³⁸ We then interact this index with year effects. We also interact the base-year share of high-skilled employment in each plant with year fixed effects and include both control variables it in $X_{f,t}$. To the extent that innovation activity and skill-biased technical change are correlated to the propensity of plants to adopting non-automation technologies, this approach should help in purging estimates of γ_1 from their impact.

Given that robot penetration in (5) varies at the 2-digit industry- and year-level, we would ideally employ industry-year two-way clustered errors. As we have only 13 industries in our data, for the baseline estimates we prefer to cluster errors at the industry-year level instead, which results in a total of 117 clusters. Section 4.5 presents robustness tests with alternative clustering (including two-way clustering).

³⁶ Cali and Presidente (2021) document significant employment spillovers to non-exposed plants from downstream automation in Indonesia.

³⁷ Similarly to Section 2.2, the downstream indicators are based on OECD average robot penetration and they are interpreted in terms of automation possibilities. This allows us to avoid concerns of endogeneity with plant-level outcomes.

³⁸ The index ranges from 0 to 1. It is computed by taking the average of five dummies, each equal to 1 if a plant had any R&D unit, performed product innovation, process innovation, used computers, or the Internet in 2006.

Finally, we normalise $secondary_{f,t_0}$ so that it has sample average value equal to zero and unitary standard deviation.³⁹ This allows us to interpret the estimated coefficient as the impact of one additional robot on a plant with high opportunities for automation.

4.4 Plant-Level Results

Table 3 displays the main results of running equation 6 on our plant data. The ETR coefficient in column (1) is positive and statistically significant. It implies that on average, one additional robot per thousand workers increases the plant's employment by 1 percent. This average plant-level estimate is consistent with the positive employment impact documented in the local labor market analysis in Section 2.2.


To test our hypothesis of decreasing returns from automation, we estimate the impact of ETR on plants with different initial exposure to robots. To that end, we first estimate an augmented version of (6), where $ETR_{f,t}$ is interacted with $dec^z(ETR_{f,t_0})$. The latter is a dummy variable identifying whether ETR_{f,t_0} is above the z th decile of its distribution. We run this regression for each decile and present the estimated coefficients and their 90 percent confidence intervals in Figure 6. There is no significant difference in robots' impact up to the median of the cumulative distribution of ETR_{f,t_0} . After that, the employment impact starts to decline gradually for plants with ETR_{f,t_0} above the 5th decile. However, the coefficients of the interaction become statistically significant only above the top two deciles, and particularly so for plants in the top decile of the distribution of initial exposure to automation.

Following this finding, we use the 9th decile threshold to identify plants with high-initial ETR .⁴⁰ We then interact this dummy with $ETR_{f,t}$ in (6). The results are presented in column (2), which shows that the employment elasticity is significantly lower for plants with high-initial exposure. This is consistent with the implications of the task-based model in Section 3.1. In particular, for high-exposed plants, one additional robot per thousand workers is associated with a positive but not statistically significant employment elasticity (coeff. = .008, s.e. =

³⁹ The actual sample average value of $secondary_{f,t_0}$ is 0.616 (see Table II).

⁴⁰ Of these, roughly 70% are in Rubber and plastics, 10% in Motor vehicles; 9% in Chemicals, and 7% in Basic metals. The rest of the high-initial exposure plants operate in Other transport equipment, Repair and installation, Electronics, and Wood and furniture.

Figure 6: Employment impact of ETR relative to less exposed plants, by decile of the distribution of initial ETR.



The figure shows point estimates and 90 percent confidence intervals of the interaction between $ETR_{f,t}$ and $dec^z(ETR_{f,t_0})$, where the latter is a categorical variable representing whether ETR_{f,t_0} is above the z th decile of its distribution. The estimates are based on model (6). Standard errors are clustered at the 2-digit industry-year-level.

Source: authors calculations based on SI and IFR.

.005).⁴¹ Instead, for the other plants, the estimated elasticity is equal to 3.3 percent (p-value = 0.008), which is significantly higher than the average effect across the entire distribution.

These results suggest that the positive employment impact documented in Section 2.2 might then be driven by plants which can exploit abundant untapped automation possibilities. In a country at the initial stage of automation, as Indonesia, such plants represent the bulk of the entire manufacturing sector population.

4.5 Robustness of Plant-Level Results

A general concern with the plant-level estimation is that since automation is concentrated in few industries (Figure 3), there might be too little variation in robot penetration to estimate the parameters of (6). However, this is unlikely to be a problem in our framework, because the identification is mainly based on *within-industry* differences in exposure to robots across plants.

⁴¹ This impact corresponds to the sum of the coefficient in column (2).

Table 3: Exposure to robots and plant-level employment: main results.

	(1) Employment	(2) Employment
ETR	0.010** (0.005)	0.033*** (0.012)
ETR \times high-initial exposure		-0.025* (0.015)
Observations	65,573	65,573
R-squared	0.927	0.927
Plant FE	yes	yes
Industry-year FE	yes	yes
Other technologies	yes	yes
Downstream automation	yes	yes

The table presents OLS estimates of the relationship between plants' exposure to robots (ETR) and log-employment. ETR is defined as industry robot penetration times the plant-level share of secondary education employment in the base year. High-initial exposure plants are plants with base-year ETR larger than the 90th percentile of the ETR distribution in the base year. Other technologies are captured by: i) an index of plant innovation activities in the base year, interacted with year fixed effects, and ii) the plant share of tertiary education workers in the base year, interacted with year fixed effects. Downstream automation captures automation possibilities in downstream industries using 5-digit \times 2-digit industry IO tables and average OECD robot penetration interacted with plant-level share of secondary education employment in the base year. Standard errors are clustered at the 2-digit industry-year-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Therefore, our identification is valid as long as there is some variation in industry penetration.

Moreover, despite the concentration in industry penetration, exposure to robots exhibits substantial variation even between industries. This is because plants in low-penetration industries might be highly exposed to robots due to a high share of secondary education workers (and vice-versa). This can be seen in Online Appendix Figure A7, which shows the industry-level average share of plants with high exposure to robots.⁴² The figure shows that all plants in Motor vehicles — the industry with the highest robot penetration — can be considered highly exposed. However, most other industries exhibit a substantial share of high-exposure plants in spite of considerably lower penetration.

Another plausible concern with the estimates of Table 3 is that they may not survive to the use of alternative specifications. We probe the results to various modifications to our specification.

First, we check if the finding of decreasing marginal returns is robust to measuring exposure to automation in the current rather than the base year. To that end, Table A9 adds ETR-squared

⁴² We deem high exposure plants those with exposure higher than the median sample value in at least one year. Results are similar with alternative definitions of high exposure.

instead of the interaction term. The estimated coefficient of the squared term is negative and statistically significant, confirming the presence of decreasing marginal returns.

Second, we address the possible concern that robot imports may be endogenous to plants' performance. We do so by instrumenting $ETR_{f,t}$ with exposure to robots based on OECD average penetration, i.e. $ETR_{f,t}^{OECD} \equiv \frac{R_{i,t}^{OECD}}{L_{i,t_0}^{OECD}} \times secondary_{f,t_0}$. Column (1) of Online Appendix Table A10 shows that the 2SLS coefficient is positive and larger than the OLS estimate.⁴³ Its magnitude implies that one robot per thousand workers increases plant-level employment by roughly 5 percent. The size of the effect is thus similar to the local labor market estimate of Section 2.2.

In column (2), we also implement the baseline test for decreasing marginal returns using the IV approach. However, including two endogenous regressors results in a weak first stage (F statistics equal to 4.724). As a result, the estimates are noisy and the interaction term albeit negative as in Table 3, it is not statistically significant. Hence, in columns (4) and (5) we split the sample between low- and high-initial exposure plants, which allows us to work with one endogenous variable at the time. The results confirm the baseline findings, with the impact of robots significantly smaller for high-initial exposure plants than for the others.

One important concern of adopting $ETR_{f,t}^{OECD}$ as an instrument is that automation possibilities might be correlated with other factors with a potential impact on Indonesian employment.⁴⁴ We perform a simple test for the validity of the exclusion restriction exploiting the fact that two industries — Textiles and Wood and paper products — do not experience any robot penetration throughout the period. If $ETR_{f,t}^{OECD}$ captures automation possibilities and relates to robot adoption only, then plant-level employment should not be affected by the instrument in a reduced form specification for plants in those industries. Column (1) of Online Appendix Table A11 presents the reduced form coefficient obtained regressing plant-level employment on $ETR_{f,t}^{OECD}$, which is positive and significant at the 1 percent level. Column (2) shows that indeed, for plants in industries with no actual Indonesian penetration, the OECD-based instrument is not statistically significant. These results suggest that $ETR_{f,t}^{OECD}$ has an impact on

⁴³The first stage exhibits an F statistics above 10. The associated first stage coefficient is positive and significant at the 99 percent (coeff. = .372, s.e. = .116).

⁴⁴ For instance, automation possibilities might be correlated to the adoption of computers, which in turn might have an impact on employment.

plant employment only through its impact on automation possibilities.⁴⁵

Finally, we test the robustness of the results to two-way clustering of the standard errors (by 2-digit industry and year). All coefficients are statistically significant at the 99 percent. Since the number of 2-digit industries and years is limited, the number of clusters available is low. Hence in Online Appendix Table A12 we report 90 percent confidence intervals and p-values obtained by implementing a wild bootstrapping procedure with 9999 repetitions.⁴⁶ Using the bootstrapping method, the coefficient in column (1) is not statistically significant at conventional levels (p-value = 0.104). However, when we include the interaction term with the dummy for high-initial exposure in column (2), both coefficients become statistically significant.

4.6 Productivity effects

The model in Section 3.1 posits that the positive employment effect of automation is a product of the positive productivity effect offsetting the negative displacement effect. This section tests the validity of this result by examining the effect of robots on plants' productivity.

To measure productivity, we estimate translog production functions using plant-level data. We exploit information on quantities and values of each product produced by plants, including multi-product ones, and construct plant-specific price indexes, as in Eslava et al. (2004) (see Online Appendix B.6). This is an advantage over contributions employing industry-level price deflators, which essentially assume that all plants within an industry face similar inputs' cost and charge the same price.⁴⁷ To account for differences in production technologies across industries, we estimate production function parameters separately for each 2-digit industry.

A key challenge in the measurement of productivity relates to the endogeneity of the firm's optimal choice of inputs, which we address by employing the control function approach of Akerberg et al. (2015). As in De Loecker and Warzynski (2012), we allow plants' exposure to

⁴⁵ The evidence supporting the validity of the plant-level instrument reinforces also our confidence in the regency-level instrument used in Section 2.2.

⁴⁶ Standard errors are informative only if the distribution of the estimated parameter is sufficiently close to a normal distribution. Since the advantage of bootstrapping is allowing to relax such an assumption, we choose to provide confidence intervals and p-values, rather than standard errors.

⁴⁷ Foster et al. (2008) discuss the bias arising when using plant revenue deflated by industry deflators. De Loecker et al. (2016) extend the analysis in the contest of unobserved variation in input prices.

robots to affect the expected value of future productivity.⁴⁸ This approach allows us to obtain a measure of quantity-total factor productivity (TFPQ), which reflects changes in pure technical efficiency, rather than in revenues. Online Appendix F describes the details of the estimation procedure.⁴⁹

We use the TFPQ measure as the dependent variable in specification 6. The results — presented in Appendix Table IV — confirm the large productivity impact of automation. One additional robot per thousand workers increases TFPQ by 7 percent on average (column 1). This effect is 10 percentage points larger for plants below the 9th decile of the initial ETR distribution, in line with the hypothesis of diminishing productivity returns to automation (column 2).

Columns (3) and (4) check the robustness of the results to using an alternative productivity variable: real marginal costs.⁵⁰ The results are qualitatively identical. The coefficient in column (3) implies that one additional robot per thousand workers lower marginal costs by 10 percent on average. This effect is again considerably larger in absolute terms in plants below the 9th decile of initial ETR distribution (column 2).

5 Robots For Economic Development

The estimates above support our hypothesis that diminishing returns to robots could explain why the employment effects of robot adoption are positive in Indonesia unlike in other countries with higher degrees of automation. A natural follow-up question is whether diminishing returns could drive differential employment impacts of robots in advanced versus developing economies. While it is beyond the scope of the paper to address this question in detail, we provide suggestive evidence consistent with the possible external validity of diminishing returns across a broad range of countries.

⁴⁸ Similarly, De Loecker et al. (2016) study the impact of trade reforms and include export dummies and import tariffs; De Loecker (2007) includes export quotas; Doraszelski and Jaumandreu (2013) include R&D expenditure, and Konings and Vanormelingen (2015) include measures of workforce training.

⁴⁹ We exclude outliers in the top and bottom three percent of the productivity distribution, which correspond to implausibly extreme values of the productivity measures.

⁵⁰ After estimating the parameters of the production functions industry by industry, we follow De Loecker and Warzynski (2012) and obtain also a measure of plants' marginal cost, based on the plants' first order conditions and flexible inputs' choice.

First, the extension of our task-based model shows that diminishing returns from automation continue to hold in general equilibrium (see Online Appendix D).⁵¹ Therefore, the model predicts that the productivity effect is stronger and employment effect more likely to be positive in countries with a lower initial level of automation.

This hypothesis is consistent with the country-industry evidence of decreasing productivity returns in [Graetz and Michaels \(2018\)](#). It is also broadly consistent with plant-level evidence from other economies with higher levels of automation. Albeit using different specifications, the estimated impact of robots on productivity in [Acemoglu et al. \(2020\)](#) for France and [Koch et al. \(2019\)](#) for Spain appear substantially lower than our estimates in both high- and low-initial penetration industries. For instance, [Koch et al. \(2019\)](#) find that robots adoption is associated to 1.3 percent increase in productivity, but only for exporting firms.⁵² [Acemoglu et al. \(2020\)](#) find a 5-year elasticity of productivity to robots adoption equal to 2.4 percent and no impact when weighting their estimates by firms' employment.⁵³

We provide further suggestive evidence based on 61 OECD and non-OECD countries and 12 industries. We restrict the period of analysis to 2007-2015, in line with the period used in sections 2 and 4. To run the analysis, we match data on employment from UNIDO with data on robots from the IFR at the industry-country-year level.⁵⁴

We regress log-employment on robot penetration, controlling for a series of fixed effects at country-industry, country-year and industry-year level. We further control for country-industry-year unobserved confounders by including a demand shifter, the average value added in other industries in the same country. We instrument robot penetration through a leave-out mean, obtained by averaging penetration in the same industry-year pair in other countries. In line with the literature and our own analysis above, this instrument should capture the component of robot imports driven by exogenous technological changes.

⁵¹ This conclusion might depend on the specific modelling choice we made to preserve analytical tractability. Generalising the finding is an interesting task left for future research.

⁵² The productivity measure in [Koch et al. \(2019\)](#) is broadly similar to our TFP, as they deflate values with firm-level price deflators and allow robots to affect expected productivity. However, unlike in this paper, they are able to observe robots adoption at the firm-level.


⁵³ The estimates in [Acemoglu et al. \(2020\)](#) are less comparable to ours. They use specification in long differences and use TFP indexes rather than applying control function methods, as we do in this paper.

⁵⁴ Specifically, we convert industries from 2-digit ISIC, which is the classification in the UNIDO data, into the broader industry classification used by the IFR (see Online Appendix B.3).

To check for diminishing returns, we flag country-industry cells with high-initial penetration as follows. For each industry, we compute the deciles of the base-year distribution of robot penetration. Then, we assign a number from 1 to 10 to countries according to the decile to which the respective country-industry cell belongs. Finally, we construct a high-exposure dummy variable equal to 1 if the country-industry cell belongs to the 10th decile, consistently with the strategy used for the Indonesian results. To account for the fact that advanced economies have on average a much higher number of robots than developing countries, we compute the high-initial penetration dummy for OECD and non-OECD members separately.

In all specifications, standard errors are clustered at the country-level. The complete estimation results are presented in Online Appendix Table [A13](#). Figure [7](#) presents point estimates and 90 percent confidence intervals for the full sample including 61 countries, and separately for OECD and non-OECD economies. We find evidence of diminishing returns in all samples considered. In the full sample, the employment impact of robots is positive for countries with low-initial penetration, but not significant on average. Instead, the impact is always positive for non-OECD economies, which include Indonesia. Consistently with our hypothesis and with the rest of the empirical literature, we find a negative and significant average effect in the most industrialized OECD economies. The average effect is driven by a large negative impact in high-initial penetration countries, which turns positive and significant when we focus on those with low-initial penetration.

Figure 7: Employment impact of robots in 61 countries and 12 industries (2007-2015).



The figure shows 2SLS estimates of the impact of robot penetration and 90 percent confidence intervals in a sample of 61 countries and 12 industries, from 2007 to 2015. The complete estimation results are presented in Online Appendix Table A13. Sources: authors' calculations based on IFR, STAN, SI.

While suggestive, these results provide some support to the possible external validity of the Indonesian plant-level results. Hence, increased robot penetration may have a more positive employment impact in developing countries than in those high-income and emerging economies the literature has focused on so far. More generally, the results of this paper cast doubt to the idea that the adoption of automation technologies in a developing country impairs its demand for labor (e.g. [Diao et al., 2021](#)). More micro-level evidence would be needed to better understand to what extent returns to robots vary across different stages of automation. To this end, it would be valuable to gather additional data on firm-level robot adoption in developing economies, particularly those at early stages of industrialisation.

6 Conclusions

This paper provided novel evidence casting doubts on the hypothesis that automation in developing countries entails a productivity/employment trade-off. The positive employment impact of robots that we observe in Indonesia — both at the local labor market and firm level — is at odds with evidence not only from advanced economies, but also from emerging ones such as

Mexico ([Artuc et al., 2019](#)) and China ([Giuntella and Wang, 2019](#)).

We explain this discrepancy by documenting the presence of diminishing productivity returns to robot adoption. As productivity drives employment demand, we argue that automation is more likely to boost employment in a country at the early stages of adoption, such as Indonesia.

In spite of the large productivity gains that we find to be induced by robots, Indonesian manufacturing still exhibits a low level of adoption. Therefore, it is crucial to better understand the potential role of barriers to adoption, and to what extent policy intervention may be appropriate to address them. Similarly, the presence of demand spillovers entailed by automation that are not internalized by firms provide a possible rationale for providing subsidies or otherwise facilitating the adoption of automation technologies in countries such as Indonesia.

More generally, the results in this paper suggest that a more systematic examination of the impact of automation in developing countries on their own economies is an important topic for future research. This should help shedding light on the extent to which automation may be an opportunity rather than a threat for their labor markets.

References

- Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo (2020). Ai and jobs: Evidence from online vacancies. Technical report, National Bureau of Economic Research.
- Acemoglu, D., C. Lelarge, and P. Restrepo (2020). Competing with robots: Firm-level evidence from france. In *AEA Papers and Proceedings*, Volume 110, pp. 383–88.
- Acemoglu, D. and P. Restrepo (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review* 108(6), 1488–1542.
- Acemoglu, D. and P. Restrepo (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives* 33(2), 3–30.
- Acemoglu, D. and P. Restrepo (2020). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy* 128(6), 2188–2244.
- Ackerberg, D. A., K. Caves, and G. Frazer (2015). Identification properties of recent production function estimators. *Econometrica* 83(6), 2411–2451.
- Adao, R., M. Kolesár, and E. Morales (2019). Shift-share designs: Theory and inference. *The Quarterly Journal of Economics* 134(4), 1949–2010.
- Adão, R., M. Kolesár, E. Morales, and X. Zhang (2020). Reg_ss: Stata module to compute confidence intervals, standard errors, and p-values in a linear regression in which the regressor of interest has a shift-share structure.
- Amiti, M. and J. Konings (2007). Trade liberalization, intermediate inputs, and productivity: Evidence from indonesia. *American Economic Review* 97(5), 1611–1638.
- Artuc, E., P. Bastos, and B. Rijkers (2018). Robots, tasks and trade.
- Artuc, E., L. Christiaensen, and H. J. Winkler (2019). *Does Automation in Rich Countries Hurt Developing Ones?: Evidence from the US and Mexico*. The World Bank.

- Autor, D. H., F. Levy, and R. J. Murnane (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics* 118(4), 1279–1333.
- Bartik (1991). Who benefits from state and local economic development policies?
- Borusyak, K., P. Hull, and X. Jaravel (2018). Quasi-experimental shift-share research designs. Technical report, National Bureau of Economic Research.
- Calì, M. and G. Presidente (2021). Automation and manufacturing performance in a developing country.
- Dauth, W., S. Findeisen, J. Suedekum, and N. Woessner (2019). The adjustment of labor markets to robots. *Journal of the European Economic Association*.
- David, H. and D. Dorn (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American economic review* 103(5), 1553–97.
- De Loecker, J. (2007). Do exports generate higher productivity? evidence from slovenia. *Journal of international economics* 73(1), 69–98.
- De Loecker, J., P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik (2016). Prices, markups, and trade reform. *Econometrica* 84(2), 445–510.
- De Loecker, J. and P. T. Scott (2016). Estimating market power evidence from the us brewing industry. Technical report, National Bureau of Economic Research.
- De Loecker, J. and F. Warzynski (2012). Markups and firm-level export status. *American Economic Review* 102(6), 2437–71.
- Diao, X., M. Ellis, M. S. McMillan, and D. Rodrik (2021). Africa’s manufacturing puzzle: Evidence from tanzanian and ethiopian firms. Technical report, National Bureau of Economic Research.
- Doraszelski, U. and J. Jaumandreu (2013). R&d and productivity: Estimating endogenous productivity. *Review of Economic Studies* 80(4), 1338–1383.

- Eslava, M., J. Haltiwanger, A. Kugler, and M. Kugler (2004). The effects of structural reforms on productivity and profitability enhancing reallocation: evidence from colombia. *Journal of development Economics* 75(2), 333–371.
- Faber, M. (2018). Robots and reshoring: Evidence from mexican local labor markets.
- World Bank (2020). Indonesia’s critical occupations list 2018 : Occupation profiles.
- Foster, L., J. Haltiwanger, and C. Syverson (2008). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review* 98(1), 394–425.
- Frey, C. B. and M. A. Osborne (2017). The future of employment: How susceptible are jobs to computerisation? *Technological forecasting and social change* 114, 254–280.
- Giuntella, O. and T. Wang (2019). Is an army of robots marching on chinese jobs?
- Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020). Bartik instruments: What, when, why, and how. *American Economic Review* 110(8), 2586–2624.
- Goos, M. and A. Manning (2007). Lousy and lovely jobs: The rising polarization of work in britain. *The review of economics and statistics* 89(1), 118–133.
- Goos, M., A. Manning, and A. Salomons (2009). Job polarization in europe. *American economic review* 99(2), 58–63.
- Goos, M., A. Manning, and A. Salomons (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American economic review* 104(8), 2509–26.
- Graetz, G. and G. Michaels (2018). Robots at work. *Review of Economics and Statistics* 100(5), 753–768.
- Hallward-Driemeier, M. and G. Nayyar (2017). *Trouble in the Making? The Future of Manufacturing-Led Development*. The World Bank.
- Humlum, A. (2019). Robot adoption and labor market dynamics. *Princeton University*.
- Javorcik, B. and S. Poelhekke (2017). Former foreign affiliates: Cast out and outperformed? *Journal of the European Economic Association* 15(3), 501–539.

- Koch, M., I. Manuylov, and M. Smolka (2019). Robots and firms.
- Konings, J. and S. Vanormelingen (2015). The impact of training on productivity and wages: firm-level evidence. *Review of Economics and Statistics* 97(2), 485–497.
- Kugler, A. D., M. Kugler, L. Ripani, and R. Rodrigo (2020). Us robots and their impacts in the tropics: Evidence from colombian labor markets. Technical report, National Bureau of Economic Research.
- Mertens, M. (2019). Micro-mechanisms behind declining labour shares: Market power, production processes, and global competition. Technical report, IWH-CompNet Discussion Papers.
- Miguel, E. and M. Kremer (2004). Worms: identifying impacts on education and health in the presence of treatment externalities. *Econometrica* 72(1), 159–217.
- Rodrik, D. (2016). Premature deindustrialization. *Journal of economic growth* 21(1), 1–33.
- Roodman, D., M. Ø. Nielsen, J. G. MacKinnon, and M. D. Webb (2019). Fast and wild: Bootstrap inference in stata using boottest. *The Stata Journal* 19(1), 4–60.
- Verhoogen, E. (2020). Firm-level upgrading in developing countries. *CDEP-CGEG Working Paper* 83.

Appendix

Table I: Summary statistics of regency-level variables involved in the labor market-level analysis.

	(1) N	(2) mean	(3) sd	(4) min	(5) max
Regency Δ ETR	284	0.165	0.316	-0.278	2.243
Regency Δ ETR (instrument)	284	0.504	0.524	-3.162	1.990
Change in total employment (log)	284	0.122	0.128	-0.178	0.615
Change in manufacturing employment (log)	284	0.121	0.376	-2.839	1.325
Change in services employment (log)	284	0.276	0.210	-0.176	1.373
Change in agriculture and mining employment (log)	284	-0.0920	0.273	-1.400	0.869
Change in RoW export by regency (log)	276	0.0771	0.380	-2.913	0.829
Population in 2007	281	793,849	704,630	29,682	5.756e+06
Share of workers with tertiary education in 2007	281	0.0482	0.0324	0.00898	0.206
Share of workers with no education in 2007	281	0.113	0.0459	0.0178	0.270
Natural resources share of output in 2007	281	0.495	0.233	0.0392	0.983
Share of employment at risk of computerisation	284	0.0138	0.0110	0	0.0856
Change in real capital stock (log)	277	1.188	0.475	-0.594	2.536
GDP per capita in 2007 (log)	281	2.934	0.638	1.757	5.518

Table II: Summary statistics of the variables involved in the plant-level analysis.

	(1) N	(2) mean	(3) sd	(4) min	(5) max
Employment (log)	66,877	3.969	0.967	2.197	9.458
TFPQ (log)	66,877	5.646	20.20	-111.4	207.2
Real marginal cost (log)	59,538	2.204	2.090	-19.16	15.02
Real investment in machinery and equipment (log)	66,877	1.761	3.599	0	19.64
Industry number of robots (1000s of workers)	66,877	0.163	0.728	0	14.08
Innovation-intensity (index)	66,877	0.216	0.275	0	1
Share of secondary education workers	66,877	0.616	0.361	0	1
Share of primary education workers	66,877	0.371	0.368	0	1
Share of tertiary education workers	66,877	0.0131	0.0440	0	1

Table III: Correlations between plant level exposure to robots and plant level investment.

	(1) Investment	(2) Investment	(3) Investment
ETR	0.113** (0.047)		
ETR (primary)		-0.105** (0.040)	
ETR (tertiary)			-0.018 (0.046)
Observations	65,573	65,573	65,573
R-squared	0.638	0.638	0.638
Plant FE	yes	yes	yes
Industry-year FE	yes	yes	yes
Other technologies	yes	yes	yes
Downstream spillovers	yes	yes	yes

The table presents OLS estimates of the relationship between plants' exposure to robots (ETR) and investment in machinery and equipment. Investment is expressed in log-transformed units: $\ln[x + (x^2 + 1)^{.5}]$. In column (1), ETR is defined as industry robot penetration times the plant-level share of secondary education employment in the base year. In column (2), ETR (primary) is defined as industry robot penetration times the plant-level share of primary education employment in the base year. In column (3), ETR (tertiary) is defined as industry robot penetration times the plant-level share of tertiary education employment in the base year. All education shares are normalised to have zero mean and unitary standard deviation. Other technologies are captured by: i) an index of plant innovation activities in the base year, interacted with year fixed effects, and ii) the plant share of tertiary education workers in the base year, interacted with year fixed effects. Downstream automation captures automation possibilities in downstream industries using 5-digit \times 2-digit industry IO tables and average OECD robot penetration interacted with plant-level share of secondary education employment in the base year. Standard errors are clustered at the 2-digit industry-year-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table IV: Exposure to robots and plant-level productivity.


	(1) TFP	(2) TFP	(3) Marginal cost	(4) Marginal cost
ETR	0.069* (0.038)	0.168*** (0.041)	-0.101 *** (0.021)	-0.283*** (0.084)
ETR \times high-initial exposure		-0.112** (0.049)		0.200** (0.087)
Observations	62,066	62,066	54,683	54,683
R-squared	0.994	0.994	0.680	0.681
Plant FE	yes	yes	yes	yes
Industry-year FE	yes	yes	yes	yes
Other technologies	yes	yes	yes	yes
Downstream automation	yes	yes	yes	yes

The table presents OLS estimates of the relationship between plants' exposure to robots (ETR), TFPQ and real marginal costs. ETR is defined as industry robot penetration times the plant-level share of secondary education employment in the base year. High-initial exposure plants are plants with base-year ETR larger than the 90th percentile of the ETR distribution in the base year. Other technologies are captured by: i) an index of plant innovation activities in the base year, interacted with year fixed effects, and ii) the plant share of tertiary education workers in the base year, interacted with year fixed effects. Downstream automation captures automation possibilities in downstream industries using 5-digit \times 2-digit industry IO tables and average OECD robot penetration interacted with plant-level share of secondary education employment in the base year. Standard errors are clustered at the 2-digit industry-year-level. Outliers in the top and bottom three percent of the productivity distribution are dropped prior to estimation. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Online Appendix (not for publication)

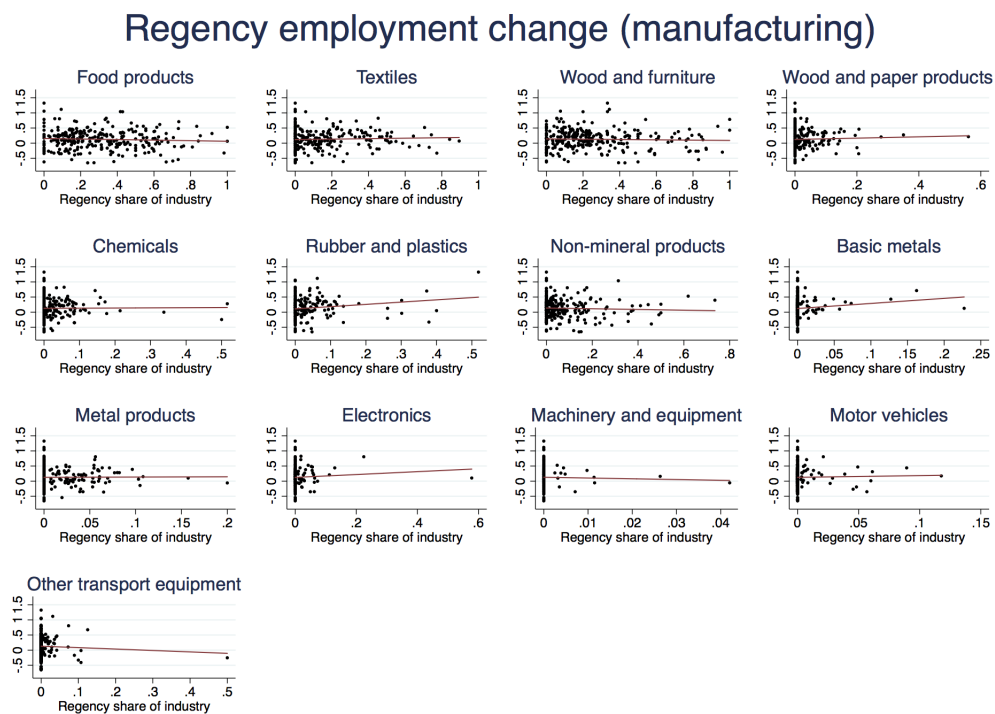
A Figures and Tables Appendix

Figure A1: Correlation between Indonesian and OECD-average exposure to robots excluding Motor Vehicles and Rubber and plastics.



On the horizontal axis there is the change between 2008 and 2015, of the OECD region industry-average number of robots per thousand workers. On the vertical axis, there is the change between 2008 and 2015, of the industry-level number of robots per thousand workers in Indonesia. The figure excludes two high-exposure industries in Indonesia, Motor vehicles and Rubber and Plastics. Sources: IFR, STAN, SI.

Figure A2: Raw correlations between base year regency-level employment shares in manufacturing industries and 2008-2015 changes in manufacturing employment.



The figure shows the correlation between the base year value of regency-level employment shares in manufacturing industries and 2008-2015 changes in manufacturing employment.

Table A1: Relationship between industry shares, ETR and regency-level characteristics.

	(1) Food products	(2) Metal products	(3) Regency ETR (instrument)
Manufacturing empl-pop ratio	2.559 (1.843)	-0.243 (0.303)	-1.766 (4.175)
GDP share of natural resources	1.894** (0.798)	-0.178 (0.144)	-0.286 (2.045)
Share of employment with tertiary education	-1.402* (0.764)	-0.084 (0.165)	0.600 (3.609)
Share of employment with no education	-0.114 (0.442)	-0.058 (0.062)	-0.395 (0.988)
Government spending on infrastructure	-0.023 (0.017)	0.004** (0.002)	-0.086* (0.051)
Real GDP per capita	-0.009 (0.029)	-0.003 (0.004)	0.106 (0.069)
GDP share of services	2.047** (0.838)	-0.111 (0.153)	-0.338 (2.216)
Population	0.033 (0.025)	0.005* (0.003)	0.155** (0.069)
Observations	275	275	275
R-squared	0.442	0.390	0.260
Province FE	yes	yes	yes

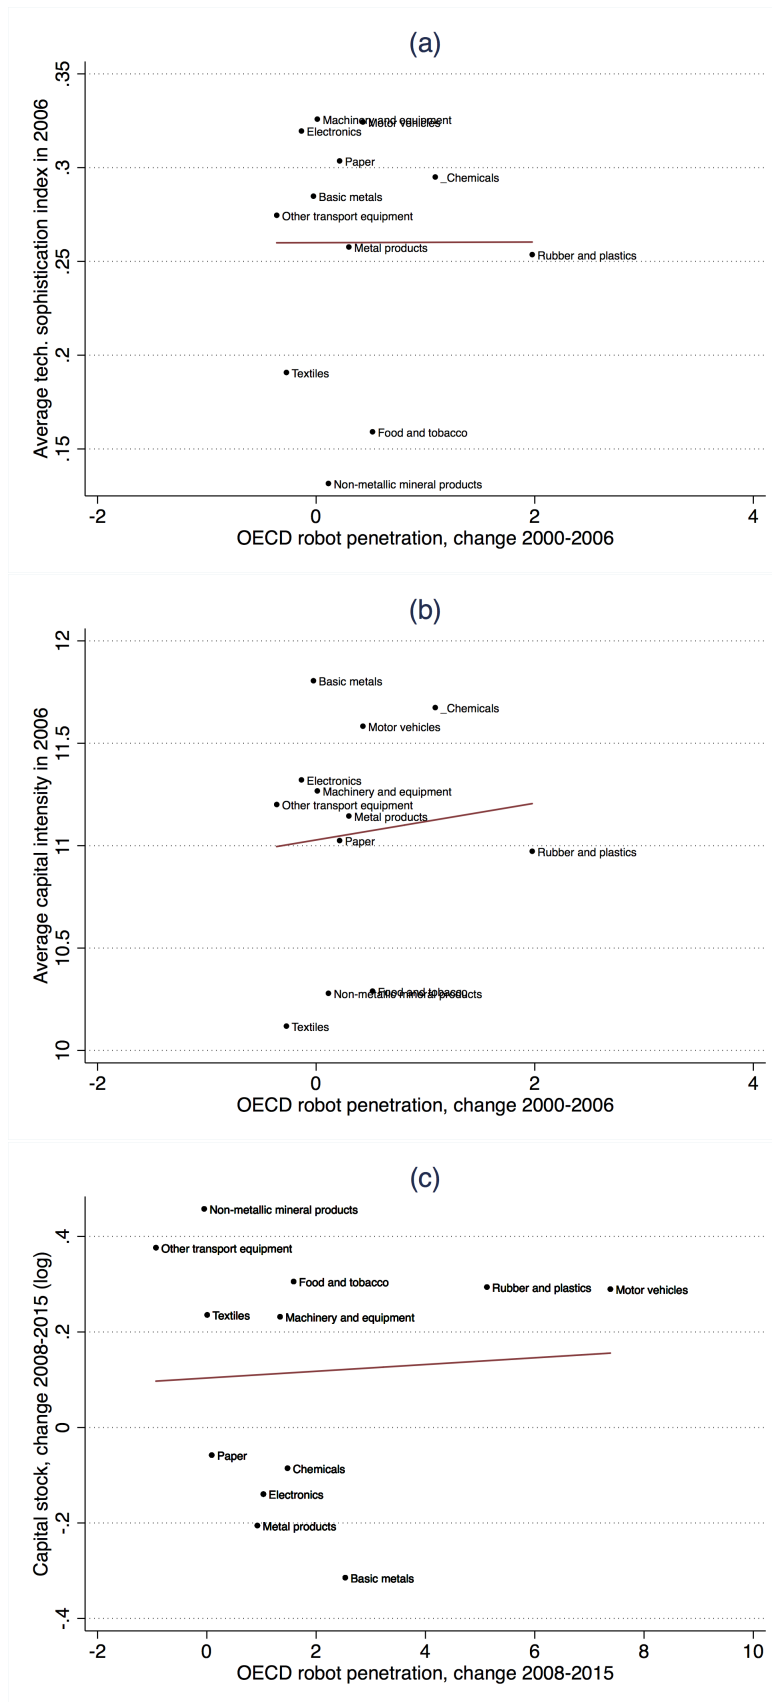
The table presents OLS estimates of the relationship between regency-level industry shares, ETR and base-year regency covariates. Industries correspond to those with the largest Rotemberg Weight. Estimates are weighted by base-year regency population. Standard errors are clustered at the regency-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table A2: Regency-level results: employment in manufacturing and non-manufacturing sectors.

	(1) Regency ETR	(2) Δ Total employment	(3) Δ Manufacturing employment	(4) Δ Services employment	(5) Δ Agr/Mining employment
Regency ETR (instrument)	0.439*** (0.078)				
Regency ETR		0.013 (0.044)	0.311*** (0.117)	0.032 (0.058)	-0.338*** (0.119)
Observations	270	270	270	270	270
R-squared	0.566				
Province FE	yes	yes	yes	yes	yes
Regency demand shifter	yes	yes	yes	yes	yes
Regency base year covariates	yes	yes	yes	yes	yes
First stage F-stat	–	31.32	31.32	31.32	31.32

The table presents 2SLS estimates of the relationship between regency-level exposure to robots and employment. Indonesian exposure to robots is instrumented with exposure based on OECD industry penetration. The dependent variables are the 2008-2015 differences of log of employment in each regency. The regency demand shifter aggregates global exports by industry (excluding Indonesia) using regency-level industry employment shares in the base year. Base-year covariates include the share of workers with no education and tertiary education (separately); the share of agriculture and mining GDP; the GDP share of services; log-real government expenditure in infrastructure, and log-population. Standard errors are clustered at the regency-level. Weights are constructed using 2007 (base year) regency population. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Figure A3: Industry-level correlations: OECD penetration and technology.



The figure presents raw correlations between industry-level 2000-2006 change in OECD robot penetration, technological sophistication (Panel (a)), and capital intensity (Panel (b)). Technological sophistication is an index based on plant-level information on IT use, R&D, and product and process innovation. Capital intensity is measured as the log of the capital-labor ratio. Panel (c) shows the correlation between industry-level 2008-2015 change in OECD robot penetration and industry-level 2008-2015 change of log-capital stock. Technological sophistication, capital intensity and the capital stock are based on plant-level data and averaged over industries.

Table A3: OECD robot penetration, aggregate capital and global shocks.

	(1) Δ Capital-labor ratio	(2) Δ Capital
Regency ETR (instrument)	-0.067 (0.089)	0.064 (0.072)
Observations	270	270
R-squared	0.275	0.369
Province FE	yes	yes
Regency demand shifter	yes	yes

The table presents OLS estimates of the relationship between regency-level exposure to robots and capital. Standard errors are clustered at the regency-level. Weights are constructed using base year regency population. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table A4: Past employment trends, ETR and industrial composition.

	(1) Regency ETR (2008-2015)	(2) Food products industry share (2007)	(3) Basic metals industry share (2007)
Manuf. employment (2000-2007)	0.016 (0.040)	0.007 (0.025)	-0.002 (0.003)
Observations	263	263	263
R-squared	0.212	0.391	0.171
Province FE	yes	yes	yes


The table presents correlations between 2000-2007 employment trends, ETR and regency-level industry shares. Standard errors are clustered at the regency-level. Weights are constructed using base year regency population. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table A5: Reduced form regency-level results

	(1) Δ Total employment	(2) Δ Manufacturing employment	(3) Δ Services employment	(4) Δ Agr/Mining employment
Regency ETR (instrument)	0.024 (0.017)	0.136*** (0.048)	0.010 (0.025)	-0.122* (0.067)
Observations	276	276	276	276
R-squared	0.012	0.039	0.002	0.053
Province FE	yes	yes	yes	yes
Regency demand shifter	yes	yes	yes	yes

The table presents OLS estimates of the relationship between regency-level exposure to robots and employment. The dependent variables are the 2008-2015 differences in log of employment in each regency. Standard errors are clustered at the regency-level. Weights are constructed using base year regency population. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Figure A4: Regency-level pre-trends: correlation between 2000-2007 and 2008-2015 changes in employment.



The figure presents raw correlations between regency-level 2000-2007 and 2008-2015 log-changes in manufacturing employment.

Table A6: Employment in manufacturing and non-manufacturing sectors: stacked differences (2008-2009/2010-2015)

	(1) Regency ETR	(2) Δ Total employment	(3) Δ Manufacturing employment	(4) Δ Services employment	(5) Δ Agr/Mining employment
Regency ETR (instrument)	0.476*** (0.076)				
Regency ETR		0.059 (0.039)	0.266*** (0.089)	-0.014 (0.052)	-0.258** (0.114)
Observations	557	557	557	557	557
R-squared	0.561				
Province FE	yes	yes	yes	yes	yes
Period FE	yes	yes	yes	yes	yes
Regency demand shifter	yes	yes	yes	yes	yes
First stage F-stat	—	38.89	38.89	38.89	38.89

The table presents 2SLS estimates of the relationship between regency-level exposure to robots and employment. Exposure to robots is instrumented with the average exposure in the OECD region. The dependent variables are the 2008-2009 and 2010-2015 stacked-differences of log of employment in each regency. The regency demand shifter aggregates global exports by industry (excluding Indonesia) using regency-level industry employment shares in the base year. Standard errors are clustered at the regency-level. Weights are constructed using 2007 (base year) regency population. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table A7: Regency-level results: other technologies.

	(1) Δ Total employment	(2) Δ Manufacturing employment	(3) Δ Services employment	(4) Δ Agr/Mining employment
Regency ETR	0.041 (0.037)	0.236** (0.107)	0.027 (0.061)	-0.270* (0.142)
Share of routine employment	0.716 (0.736)	-4.625** (2.227)	0.670 (1.152)	-0.568 (1.211)
Technological sophistication index	0.259*** (0.086)	0.140 (0.223)	0.279** (0.120)	-0.011 (0.165)
Capital-labor ratio	0.028* (0.016)	0.031 (0.043)	-0.010 (0.021)	-0.022 (0.024)
Observations	274	274	274	274
Province FE	yes	yes	yes	yes
Regency demand shifter	yes	yes	yes	yes
First stage F-stat	28.30	28.30	28.30	28.30


The table presents 2SLS estimates of the relationship between regency-level exposure to robots and employment. Exposure to robots is instrumented with the average exposure in the OECD region. The dependent variables are the 2008-2015 differences of log of employment in each regency. The share of regency employment at risk of computerisation is based on data from [Frey and Osborne \(2017\)](#). Technological sophistication is an index based on information on IT use, R&D, and product and process innovation. Capital intensity is measured as the log of the capital-labor ratio. The latter two variables are based on base-year plant-level data and averaged over regencies. The regency demand shifter aggregates global exports by industry (excluding Indonesia) using regency-level industry employment shares in the base year. Standard errors are clustered at the regency-level. Weights are constructed using 2007 (base year) regency population. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table A8: Regency-level results: alternative standard errors

	(1) Δ Total employment	(2) Δ Manufacturing employment	(3) Δ Services employment	(4) Δ Agr/Mining employment
Regency ETR	0.054 (0.069) [0.084]	0.309 (0.118) [0.103]	0.023 (0.057) [0.069]	-0.277 (0.144) [0.186]
Observations	276	276	276	276
Province FE	yes	yes	yes	yes
Regency demand shifter	yes	yes	yes	yes
Regency base year covariates	yes	yes	yes	yes


The table presents 2SLS estimates of the relationship between regency-level exposure to robots and employment. Exposure to robots is instrumented with the average exposure in the OECD region. The dependent variables are the 2008-2015 differences in log of employment and wages (total or by sector) in each regency. Standard errors in round brackets are clustered at the province-level. Standard errors in square brackets are computed with the methodology of [Adao et al. \(2019\)](#).

Figure A5: Correlation between the 2006 shares of workers with secondary education and robots in use in 2015, by industry.




On the horizontal axis there is the industry-average share of workers with secondary education, computed from plant level data in the base year. On the vertical axis there is, for each industry, the log number of robots per thousand workers in 2015. Sources: IFR, SI.

Figure A6: Correlation between IFR data on robots and SI data on investment, by industry and year.



The figure shows the correlation between industry-level IFR data on robots and plant level investment in machinery and equipment. On the horizontal axis there is the number of industrial robots shipped to each manufacturing industry and year in Indonesia (dropping observations with zero value). On the vertical axis there is the industry-year-average of plant level investment in machinery and equipment. Sources: IFR, SI.

Figure A7: Industry shares of plants with ETR higher than the sample median value in at least one year.



The figure presents the average industry share of plants with exposure to robots (ETR) larger than the sample median value in at least one year of the sample.
Source: authors calculations based on SI and IFR.

Table A9: Current-period diminishing returns.

	(1) Employment
ETR	0.011** (0.005)
ETR-squared	-0.001* (0.000)
Observations	65,573
R-squared	0.927
Plant FE	yes
Industry-year FE	yes
Other technologies	yes
Downstream automation	yes

The table presents OLS estimates of the relationship between plants' exposure to robots (ETR) and log-employment. ETR is defined as industry robot penetration times the plant-level share of secondary education employment in the base year. Other technologies are captured by: i) an index of plant innovation activities in the base year, interacted with year fixed effects, and ii) the plant share of tertiary education workers in the base year, interacted with year fixed effects. Downstream automation captures automation possibilities in downstream industries using 5-digit \times 2-digit industry IO tables and average OECD robot penetration interacted with plant-level share of secondary education employment in the base year. Standard errors are clustered at the 2-digit industry-year-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table A10: Exposure to robots and plant-level employment: 2SLS estimates.

	(1) Employment	(2) Employment	(3) Employment low exposure	(4) Employment high exposure
ETR	0.047* (0.025)	0.188* (0.107)	0.427* (0.224)	0.018** (0.008)
ETR \times high-initial exposure		-0.165 (0.102)		
Observations	53,194	53,194	48,240	4,954
Plant FE	yes	yes	yes	yes
Industry-year FE	yes	yes	yes	yes
Other technologies	yes	yes	yes	yes
Downstream automation	yes	yes	yes	yes
First stage F-stat	10.31	4.724	64.59	38.12

The table presents 2SLS estimates of the relationship between plants' exposure to robots (ETR) and employment. ETR is defined as industry robot penetration times the plant-level share of secondary education employment in the base year. ETR (OECD) is defined as OECD average industry robot penetration times the plant-level share of secondary education employment in the base year. High-initial exposure plants are plants with base-year ETR larger than the 90th percentile of the ETR distribution in the base year. Other technologies are captured by: i) an index of plant innovation activities in the base year, interacted with year fixed effects, and ii) the plant share of tertiary education workers in the base year, interacted with year fixed effects. Downstream automation captures automation possibilities in downstream industries using 5-digit \times 2-digit industry IO tables and average OECD robot penetration interacted with plant-level share of secondary education employment in the base year. Standard errors are clustered at the 2-digit industry-year-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table A11: OECD average exposure to robots: testing the validity of the exclusion restrictions.

	(1)	(2)
	Employment	Employment (no penetration)
ETR (OECD)	0.018*** (0.006)	0.271 (0.224)
Observations	53,194	16,923
R-squared	0.929	0.929
Plant FE	yes	yes
Industry-year FE	yes	yes
Other technologies	yes	yes
Downstream automation	yes	yes

The table presents the reduced form impact of ETR (OECD) on plant-level employment. ETR (OECD) is defined as OECD average industry robot penetration times the plant-level share of secondary education employment in the base year. Column (1) presents the average impact on the full sample. Column (2) includes only Textile and Wood and paper products, which experience no robot penetration in Indonesia over the years of the sample. Other technologies are captured by: i) an index of plant innovation activities in the base year, interacted with year fixed effects, and ii) the plant share of tertiary education workers in the base year, interacted with year fixed effects. Downstream automation captures automation possibilities in downstream industries using 5-digit \times 2-digit industry IO tables and average OECD robot penetration interacted with plant-level share of secondary education employment in the base year. Standard errors are clustered at the 2-digit industry-year-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table A12: Exposure to robots and plant-level employment: alternative two-way clustering and (two-way) wild bootstrapping.

	(1) Employment	(2) Employment
ETR	0.010*** (0.002) [-.0002, .0186] {0.104}	0.033*** (0.007) [.0120, .0541] {0.048}
ETR × high-initial exposure		-0.025*** (0.005) [-.0512, -.0015] {0.088}
Observations	65,573	65,573
R-squared	0.927	0.927
Plant FE	yes	yes
Industry-year FE	yes	yes
Other technologies	yes	yes
Downstream automation	yes	yes

The table presents OLS estimates of the relationship between plants' exposure to robots (ETR) and employment. ETR is defined as industry robot penetration times the plant-level share of secondary education employment in the base year. High-initial exposure plants are plants with base-year ETR larger than the 90th percentile of the ETR distribution in the base year. Other technologies are captured by: i) an index of plant innovation activities in the base year, interacted with year fixed effects, and ii) the plant share of tertiary education workers in the base year, interacted with year fixed effects. Downstream automation captures automation possibilities in downstream industries using 5-digit × 2-digit industry IO tables and average OECD robot penetration interacted with plant-level share of secondary education employment in the base year. Standard errors in parenthesis are tow-way-clustered at the 2-digit industry and year-level. Square brackets include 90% confidence intervals based on two-way a wild bootstrapping procedure (over 9999 repetitions) implemented with the user-written STATA command *boottest* (Roodman et al., 2019). Curly brackets present the associated p-values.

Table A13: Employment impact of robots in 61 countries and 12 industries (2007-2015).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	All countries			Non-OECD			OECD		
	All industries	High penetration	Low penetration	All industries	High penetration	Low penetration	All industries	High penetration	Low penetration
Robot penetration	-0.004 (0.008)	-0.007 (0.006)	0.118*** (0.017)	0.003* (0.001)	0.010*** (0.003)	0.095* (0.054)	-0.025*** (0.010)	-0.057*** (0.001)	0.106*** (0.013)
Observations	5,447	358	4,864	2,856	144	2,484	2,591	142	2,380
R-squared	0.077	0.456	0.089	0.113	0.769	0.108	0.058	0.214	0.078
Demand shifter	yes	yes	yes	yes	yes	yes	yes	yes	yes
Country-industry FE	yes	yes	yes	yes	yes	yes	yes	yes	yes
Country-year FE	yes	yes	yes	yes	yes	yes	yes	yes	yes
Industry-year FE	yes	yes	yes	yes	yes	yes	yes	yes	yes

The table presents 2SLS estimates of the relationship between log-employment and robot penetration. We instrument robot penetration with a leave-out mean, obtained by averaging penetration in the same industry-year pair in other countries. High penetration industries are those with base-year penetration above the 9th decile of the initial penetration distribution computed across countries (for OECD and non-OECD separately). The demand shifter is the average value added in other industries in the same country. Standard errors are clustered at the country-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

B Data Appendix

B.1 Robots Data

One issue with IFR data is that in the early years of the sample, a breakdown of imports by sector is not available and they are grouped under the label “unspecified”. In this case, shares by sectors are estimated using information for the years in which the breakdown is available. We experiment with two alternatives, namely taking simple averages over all the available years and using the observation for the most recent available year. Results are very similar. The resulting shares are used to construct the deliveries by sector. As in [Graetz and Michaels \(2018\)](#), the construction of the stock of operational robots is obtained by assuming a yearly depreciation rate of 10% and applying the perpetual inventory method, using 1993 estimates of the existing stock by the IFR as initial values. The IFR does provide estimates of the stock, but it adopts a different assumption that robots fully depreciate after twelve years.

The original IFR industry classification has been converted to obtain eighteen industries, roughly corresponding to 2 digit-level ISIC rev.4. These are: Agriculture, Food and tobacco, Textiles, Paper, Wood and furniture, Chemicals, Rubber and plastics, Non-metallic mineral products, Basic metals, Metal products, Electronics, Machinery and equipment, Motor vehicles, Other transport equipment, Repair and installation of machinery, Construction, and Education and R&D, and Utilities.

B.2 Manufacturing Data

The SI survey is administered by the Indonesian statistical office (BPS) and its coverage is extensive. In fact it becomes an actual census in 2006 and it is very close to a census in the remaining years, hence ensuring high representativeness even at very low levels of aggregation. Importantly, the 2006 Indonesian census includes also plant level information on employment by educational attainment, which allows to construct the plant-based measures of exposure to robots as in the main text. In particular, the educational categories in our plant level data are: i) Not finished primary school; ii) Primary school; iii) Junior high school; iv) Senior high school; v) Diploma; vi) Bachelor; vii) Master, and viii) Ph.D. We define as “secondary” educational

attainments falling into the categories iii) and iv).

In addition to plant level output, capital, labor and educational attainments, SI data provide information also on the quantity and value of 9-digit products produced and input used (domestically produced and imported) by each plant.⁵⁵

One challenge of the Statistik Industri data is the lack of complete series of capital stock. Earlier studies tried to re-construct capital stock series applying the perpetual inventory method (PIM) to the first year of capital stock data reported by the plant (Amiti and Konings (2007); Javorcik and Poelhekke (2017)). However this imputation method crucially relies on the capital value self-reported by the plant the first year this data is available, which is not necessarily accurate.⁵⁶ One potential advantage of using PIM is that purchase and sales data might be more accurate relative to self-reported value of the stock, requiring an appropriate calculation of market values and depreciation. However, PIM needs to rely on measures of capital depreciation, which are difficult to accurately estimate. To mitigate such trade-off, we have adopted a hybrid strategy. We first clean the self-reported adopting an algorithm which keeps only observations that fulfill a battery of tests, which are described in Appendix B.4. Then, we apply the PIM only to fill the gaps between the missing observations and reapply the same battery of tests to ensure consistency of the series.

In order to allow the matching between SI and IFR data, we build a consistent industry classification. Plants in SI are grouped into 5-digit sectors following the definition Klasifikasi Baku Lapangan Usaha Indonesia (KBLI). A KBLI code is assigned to a plant according to the classification in which the main product produced belongs. The KBLI classification has been adjusted to be consistent over the whole sample, ranging from 2006 to 2015. One issue is that in converting codes from KBLI rev.3 (in use until 2009) to KBLI rev.4, some industries are split in more than one industry, or vice-versa. For such reason, we only keep those KBLI codes that have an unambiguous one to one mapping across the two revisions. We also experimented with a looser conversion including more industries, without significant changes in our main results.

⁵⁵ In our sample, each plant produces on average 2 products and 25% of the plants produce more than one product. On average, each plant uses four different varieties of raw inputs.

⁵⁶ In particular, there is no a priori reason to believe that the quality of the self-reported capital stock the first year is necessarily better than the value in other years.

B.3 Matching SI and IFR Data

The KBLI classification of SI data is very similar to the ISIC Rev. 4 coding of the IFR data. However, in some cases the SI data are more detailed than the IFR ones. Thus, we group together some KBLI industries to ensure maximum compatibility across the two datasets. The correspondence is shown in Table B1. We observe that 8.5% of plants switches to another industry during the 10 years covered by our sample. Therefore, to avoid potentially confounding effects, we assign to each plant the trends in robot adoption of the industry to which it belonged in the first year of observation.

Table B1: Correspondence between IFR and SI industry classification.

IFR industries	Description	KBLI industries
D10T12	Food products, beverages and tobacco	10,11,12
D13T15	Textiles, wearing apparel, leather and related products	13,14,15
D16and3132	Wood, furniture, n.e.c	16,31,32
D17T18	Wood and paper products	17,18
D19T21	Chemicals	19,20,21
D22	Rubber and plastics products	22
D23	Other non-mineral products	23
D24	Basic metals	24
D25	Metal products	25
D26T27	Electronics	26,27
D28	Machinery and equipment n.e.c.	28
D29	Motor vehicles	29
D30	Other transport equipment	30

B.4 Construction of the Capital Series

In order to avoid relying on depreciation rates, we tried to preserve the self-reported original values by the plant as much as possible and applied the PIM only to fill gaps. In this paper self-reported capital series were object of an extensive cleaning algorithm aimed at mitigating measurement errors. One problem with the reported series is that in some years, there are plants were characterised by implausible large values of capital. Studying the behavior of the stock within plants reveals that in some circumstances plants reported values in different units. The phenomenon is somewhat more frequent in 1996 and 2006, when the BPS conducted a wider economic census that collected information in units rather than in thousand Rupiah. For instance, in 2006 the number of surveyed firms increased by 40%. The increase in coverage required hiring inexperienced enumerators that were more likely to make mistakes, which contributed to increase measurement errors.

Our algorithm consists first in replacing zero or negative values as missing observations and then applying a two-steps procedure based on capital-labor ratios (KL). For each year, we compute the average KL in each 4-digit KBLI sector over the whole sample, but excluding the years in which the average and total values of the capital stock exhibited suspicious jumps, i.e. 1996, 2000, 2003, 2006, 2009 and 2014. An observation is dropped if the ratio of plant-KL to the sector average KL is below 0.02 or larger than 50. We experiment with stricter thresholds which result in too many observations dropped. Then, in a second step we compare a plant KL in a given year with the average value of the KL within the same plant but in the other years of observation. An observation is dropped if the ratio of plant-year-KL to the plant average KL is below 0.2 or larger than 5. Plants are dropped from the sample in case the cleaning procedure results in all missing values of self-reported capital. When a plant has some but not all valid observations for self-reported capital stock, then missing values are replaced by applying a forward/backward perpetual inventory method (PIM). Being only a fraction of the total observations, we rely less on estimates of depreciation rates. We follow Arnold and Javorcik (2009) and assume that the annual depreciation rate for buildings is 3.3 percent, for machinery 10 percent, and for vehicles and other fixed assets 20 percent. For land, we assumed no depreciation.

Previous studies focus on the first year of observation of a plant, without assessing the plausibility of the data point. Since PIM series are very sensitive to the choice of the initial observation, especially with relatively short time series, the resulting capital stock could be severely mis-measured. Moreover, information on purchases and sales of capital equipment, which is subject to the same measurement errors of the reported capital. For such a reason, after filling missing values with the PIM we re-apply the two stages check described above in order to minimize the possibility of mis-measurement. As a final test, we compute plant level growth rates of KL and we check that it is reasonably distributed (Figure B1). Figure B2 compares original and clean capital stock series.

Figure B1: Plants' growth rate distribution of capital-labor ratio.





Figure B2: Comparison of Aggregate Nominal Capital Stock Series.



B.5 Nine-digit Products and Inputs Data

Our data include information on quantities and values of the products produced and raw materials used by each plant. These are both defined at a highly granular level, namely 9-digit Klasifikasi Komoditi Indonesia (KKI). In our sample, each plant produces on average 2 products and 25% of the plants produce more than one product. We use disaggregate products information to measure the number of products produced by each plant. After computing unit prices by dividing value with quantities, we compute yearly price growth. If the price grow by more than a factor of 10 or decreases more than by a factor of 1/10, we drop the observation. Average unit prices are then used to construct plant level price deflators (see B.6). On average, each plant uses four different varieties of raw materials. We also have information on use of domestically produced and imported materials, which we aggregate at the plant level to measure the share of imported materials.

B.6 Construction of plant level Price indices

The derivation of plant-specific price indices from product-level price data closely follows [Eslava et al. \(2004\)](#) and [Mertens \(2019\)](#).

These are plant level Tornqvist indices exploiting information on 9-digit products produced and inputs used by each plant.

$$\pi_{jt} = \prod_{p=1}^n \left(\frac{P_{pjt}}{P_{pj,t-1}} \right)^{.5(s_{pjt} + s_{pj,t-1})} \pi_{j,t-1}$$

where P_{pjt} is the price of good p and s_{pjt} is the share of this good in total product market sales of plant j in period t . Therefore, the growth of π_{jt} is the product of each plant's price growth, each weighted with the average share of sales in t and $t - 1$. We set $\pi_{jt} = 100$ in 2006. For plants entering after 2006, we follow [Eslava et al. \(2004\)](#) and [Mertens \(2019\)](#) and use the 5-digit industry average of the plant price indices as a starting value. When price growth data are missing, we replace it with an average of product or inputs price changes within the same 5-digit industry.

C Occupation Profiles

The Occupation profiles are compiled by the World Bank in partnership with national governments. The Occupational profiles are used as key inputs of Critical Occupations Lists, which aim to identify shortages of certain occupations of strategic importance to the economy (e.g. [World Bank \(2020\)](#)).

The methodology used by the world bank is based on the probability of computerization provided by [Frey and Osborne \(2017\)](#). Occupations with probability greater than 0.7 are deemed at high risk of automation.

The list of occupations at high risk of automation extracted from the World Bank Occupation profiles for Indonesia, Malaysia and Thailand are: Welders and Flame Cutter; Handicraft Workers in Wood, Textile, Leather and Related Materials; Power-Production Plant Operators; Woodworking-machine tool setters and operators; Mineral and stone processing plant operators; Cement, stone and other mineral products machine operators; Well drillers and borers and related workers; Cement, stone and other mineral products machine operators; Metal processing plant operators; Metal finishing, plating and coating machine operators; Chemical products plant and machine operators; Rubber products machine operators; Plastic products machine operators; Food and related products machine operators; Packing, bottling and labelling machine operators; Stationary plant and machine operators; Mechanical machinery assemblers; Electrical and electronic equipment assemblers.

The information provided in the World Bank occupation profiles is similar to the Occupational Information Network database (O*NET) for the United States, but it is based on analyses of labor force survey data for comparable Asian economies. An example of the Occupation Profile for “Welders and Flame Cutters” is provided in Figure [C1](#). According to the World Bank, this is an occupation at high risk of automation.


The key piece of information provided in the Occupation Profiles is the typical educational attainments of workers in occupations at high risk of automation. An example of occupation profile for “Welders and Flame Cutters” is provided in Figure [C1](#). As for Welders and Flame cutters, a manual inspection of all available occupations suggests that the modal educational attainment for occupations at high risk of automation is junior or senior secondary education—

items iii) and iv) in the SI educational category (see Section ??).

Figure C1: Example of occupation profile, compiled by the World Bank.

Welders and Flame Cutter

Summary	
KBJI title:	Welders and Flame Cutter
KBJI code:	7212
Job description:	Workers in this group weld and cut metal parts using gas flames or electric flares and melt and unify the metal. Tasks include: to weld parts of metal using gas or electric flares, thermite compounds, or other methods; operate resistance welding machines; use a welding torch to make and repair important layers, pipes, floors, and other equipment; solder metal parts; cut metal plates with fire or electric incandescent; combine metal parts with hand soldering.
Position includes:	Welder (with specification on underwater welders and welders for food and beverage manufacturing)
Statistics	
Gender:	1% female
Age:	33-38 years old
Monthly wage:	IDR 1,800,000 – 2,375,000
Working hours:	47-48 hours per week
Typical education:	Junior High School and Vocational High School
Field of study:	N/A
Most common sector of employment:	Manufacturing
Most common employer type:	Individual/Household business



Sources: World Bank Occupation Profiles

D Model Appendix

D.1 Setup

The final good is produced by a representative firm combining a continuum of varieties of total measure equal to 1. Each variety is produced by an intermediate good-producing firm f :

$$Y = \left[\int_0^1 y_f^{\frac{\sigma-1}{\sigma}} df \right]^{\frac{\sigma}{\sigma-1}}$$

with $\sigma > 1$.

We use the price of the final good as the numeraire $P \equiv 1$. Thus, each firm faces a constant elasticity demand function:

$$y_f = p_f^{-\sigma}$$

The firm produces the variety combining a unit measure of tasks, each indexed by z , through the production function

$$y_f = \exp \left(\int_0^1 \ln x_f(z) dz \right)$$

where $x_f(z)$ is the quantity of task z demanded by the firm. Since all firms are identical and face the same problem, we can suppress f subscripts.

Tasks can be performed by human workers or machines with the following task production functions:

$$x(z) = \begin{cases} \gamma(z) \cdot n(z) & \text{if performed with labor} \\ \eta(z) \cdot k(z) & \text{if performed with capital} \end{cases} \quad (\text{D1})$$

where $n(z)$ and $k(z)$ are labor and capital allocated by the firm to the production of task z . We assume labor and capital to be fully flexible across tasks and firms.

The ratio $\gamma(z)/\eta(z) \equiv \tilde{\gamma}(z)$ represents the comparative advantage of human labor over machines in performing task z . Without loss of generality, we assume that $\tilde{\gamma}(z)$ is increasing

in z , so that higher-ordered tasks are harder to automate.

D.2 Producers' Partial Equilibrium

The firm-level intensive margin of automation in the model is summarised by a single parameter, κ . Tasks $z \in [0, \kappa]$ are automated, while tasks $z \in (\kappa, 1]$ are performed by labor.

Given the linearity of the task production function (D1), factor demand does not depend on factor prices. Labor and capital are allocated across tasks as follows:

$$x^*(z) = \begin{cases} \frac{n}{1-\kappa} & \text{if } z \in (\kappa, 1] \\ \frac{k}{\kappa} & \text{if } z \in [0, \kappa] \end{cases} \quad (\text{D2})$$

Plugging $x^*(z)$ into the production function, we get

$$y = \exp \left(\int_0^\kappa \ln \eta(z) dz + \int_\kappa^1 \ln \gamma(z) dz \right) \left(\frac{k}{\kappa} \right)^\kappa \left(\frac{n}{1-\kappa} \right)^{1-\kappa} \quad (\text{D3})$$

Capital and labor are perfectly mobile across firms. Firms are monopolistically competitive and solve:

$$\max_{\{k, n\}} py - rk - wn$$

s.t. (D3) and $y = p^{-\sigma}$

The first order conditions for labor and capital are given by

$$wn = (1 - \kappa) \left(1 - \frac{1}{\sigma} \right) py \quad (\text{D4})$$

$$rk = \kappa \left(1 - \frac{1}{\sigma} \right) py. \quad (\text{D5})$$

Solving for k in (D4) and (D5):

$$k = \frac{\kappa}{1 - \kappa} \frac{w}{r} n \quad (\text{D6})$$

Notice that we can also express the factor price ratios in terms of the capital-labor ratio and the share of automated tasks:

$$\frac{w}{r} = \frac{k}{n} \frac{1 - \kappa}{\kappa} \quad (\text{D7})$$

Substituting (D6) in the production function (D3) and taking logs, we get

$$\ln y = \int_0^\kappa \ln \eta(z) dz + \int_\kappa^1 \ln \gamma(z) dz + \ln n - \ln(1 - \kappa) + \kappa \ln \left(\frac{w}{r} \right) \quad (\text{D8})$$

Taking the log of (D4) and substituting (D8) into it, we obtain an expression for labor:

$$\ln n = (\sigma - 1) \left[\int_0^\kappa \ln \eta(z) dz + \int_\kappa^1 \ln \gamma(z) dz + \kappa \ln \left(\frac{w}{r} \right) \right] + \ln(1 - \kappa) + \Gamma \quad (\text{D9})$$

$$\text{with } \Gamma \equiv \ln \left[w^{-\sigma} \left(1 - \frac{1}{\sigma} \right)^\sigma \right].$$

Differentiating (D9) with respect to κ leads us to the following proposition.

Proposition 1 *Employment effect in partial equilibrium*

(a) *The employment effect of automation is given by*

$$\frac{\partial \ln n}{\partial \kappa} = \frac{\sigma - 1}{\sigma} \frac{\partial \ln y}{\partial \kappa} - \frac{1}{1 - \kappa} \quad (\text{D10})$$

The employment effect is the sum of a term proportional to the productivity effect, $\frac{\partial \ln y}{\partial \kappa}$, and a displacement effect, which is always negative, $-\frac{1}{1 - \kappa} < 0$.

We now characterise the productivity effect. Substituting back (D9) into (D8), we get

$$\ln y = \sigma \left[\int_0^\kappa \ln \eta(z) dz + \int_\kappa^1 \ln \gamma(z) dz + \kappa \ln \left(\frac{w}{r} \right) \right] + \Gamma \quad (\text{D11})$$

Differentiating (D11) with respect to κ leads us to the following proposition.

1. Proposition 2 *Productivity effect in partial equilibrium*

(a) *The productivity effect of automation is given by*

$$\frac{\partial \ln y}{\partial \kappa} = \sigma \left[\ln \left(\frac{w}{r} \right) - \tilde{\gamma}(\kappa) \right] \quad (\text{D12})$$

The productivity effect is positive as long as $\frac{w}{\gamma(\kappa)} > \frac{r}{\eta(\kappa)}$, i.e. the cost of labor in efficiency units is larger than the cost of capital in efficiency units.

(b) *Automation has diminishing productivity returns:*

$$\frac{\partial^2 \ln y}{\partial \kappa \partial \kappa} = -\sigma \tilde{\gamma}'(\kappa) < 0$$

Proposition 2 emphasises two aspects of the productivity effect in partial equilibrium. Part (a) of the proposition says that even in labor-abundant economies, where w/r tends to be small, automation can have a positive impact on firms' productivity if labor productivity is low. Despite having a low cost of labor, machines could still outperform labor due to their greater speed or precision. This is likely to be the case in developing countries, where labor quality tends to be low.

Part (b) of Proposition 2 shows that the productivity effect is more likely to be large with initial low levels of automation, i.e. small κ . This derives from the fact that the $\tilde{\gamma}(\cdot)$ is increasing, which reflects the comparative advantage of humans in performing non-routine, complex tasks. In a developing country, where only a relatively small share of tasks is automated, there are more opportunities for productivity gains. On the contrary, in an advanced economy, where producers are more likely to have exploited all available opportunities to improve productivity through automation, using machines to perform the marginal task would deliver limited gains due to the strong comparative advantage of humans in performing complex tasks that are not yet automated.

It should be noticed that the decreasing productivity returns in partial equilibrium arise because of our assumption on labor comparative advantage on certain tasks. If capital and labor

had the same marginal productivity on all tasks, i.e. $\tilde{\gamma}(z) = \tilde{\gamma}$, then automation would have increasing productivity returns because $\frac{\partial^2 \ln y}{\partial \kappa \partial \kappa} = 0$. This suggests that in partial equilibrium, the task-based model exhibits decreasing returns to automation for a different reason than in a standard neo-classical model, where decreasing returns to capital are solely due to the concavity of the production function.

D.3 General Equilibrium

We now turn to characterise the general equilibrium of the model.

The consumers side of the economy is composed by a representative household with quasi-linear preferences:

$$U(C, n, k) = C - \frac{n^{1+\frac{1}{\varepsilon_n}}}{1+\frac{1}{\varepsilon_n}} - \frac{k^{1+\frac{1}{\varepsilon_k}}}{1+\frac{1}{\varepsilon_k}}$$

where C denotes consumption of the final good, ε_n is the inverse labor supply elasticity, which fully parametrizes the disutility of supplying labor. In order to avoid introducing dynamics in the model, we assume that the household transforms the consumption good into capital at some cost, which is parametrized by the inverse capital supply elasticity, ε_k .

The budget constraint of the household is given by

$$C = wn + rk$$

where w is the wage rate and r the price of capital.

The optimality conditions for the representative household are given by the labor and capital supply schedules:

$$n = w^{\varepsilon_n}$$

$$k = r^{\varepsilon_k}$$

Inverting the supply curves, taking their ratio and manipulating the expression, we can

substitute the expression for w/r into (D12):

$$\frac{\partial \ln y}{\partial \kappa} = \sigma \left[\frac{\varepsilon_k - \varepsilon_n}{\varepsilon_n \varepsilon_k} \ln n - \frac{1}{\varepsilon_k} \ln \left(\frac{k}{n} \right) - \tilde{\gamma}(\kappa) \right] \quad (\text{D13})$$

Equation (D13) shows that the productivity effect is composed by three terms. If capital supply is more elastic than labor supply, the first term is negative and decreasing in employment. This is likely the case in less developed economies, where getting capital is trickier than in advanced ones. Increasing automation would be detrimental to productivity because it is cheaper the alternative of using abundant labor. This term thus expresses the idea that developing countries should not automate in order to exploit their comparative advantage.

However, the second term in (D13), a function of the capital-labor ratio, has negative sign. In situations where $k < n$, as it might be the case in less developed economies, this term is positive. Therefore, automation should be particularly beneficial for under-capitalized firms. Moreover, if capital supply is very elastic, $\varepsilon_k \simeq 0$ and so the second term contributes substantially to determining the value of $\frac{\partial \ln y}{\partial \kappa}$.

Finally, the last term in equation (D13) represents the comparative advantage of labor in complex tasks, which is decreasing in the initial level of automation, as discussed in Section D.2.

We now fully solve the model to express the productivity effect in terms of κ and model's parameters. To do so, we need the ideal price condition:

$$P \equiv 1 = \int_0^1 \left(r^\kappa w^{1-\kappa} \right)^{1-\sigma} df \quad (\text{D14})$$

Solving the system of equations, we obtain an expression for the equilibrium factor price ratio:

$$\ln \frac{w}{r} = \frac{1 - \kappa}{\kappa} \frac{\ln(1 - \kappa) - \ln \kappa}{(1 - \kappa)(1 + \varepsilon_n) + \kappa(1 + \varepsilon_k)} \quad (\text{D15})$$

Substituting (D15) into (D12) leads us to the following proposition.

1. Proposition 3 *Productivity effect in general equilibrium*

The productivity effect of automation is given by

$$\frac{\partial \ln y}{\partial \kappa} = \sigma \left[\frac{1 - \kappa}{\kappa} \frac{\ln(1 - \kappa) - \ln \kappa}{(1 - \kappa)(1 + \varepsilon_n) + \kappa(1 + \varepsilon_k)} - \tilde{\gamma}(\kappa) \right] \quad (\text{D16})$$

The first term in square brackets in (D16) is strictly positive and monotonically decreasing in the interval $\kappa \in \left(0, \frac{1}{2}\right)$.

Proposition 3 illustrates that in our task-based model, the productivity effect of automation is very strong for initial low levels of automation and then it decreases rapidly. This is due to the decreasing returns to capital generated by the concavity of the production function, but also to $\tilde{\gamma}(\cdot)$, which captures the comparative advantage of labor in the most complex tasks.

Finally, we can express the employment impact of automation as a function of the productivity effect and the displacement effect. Differentiating (D9) with respect to κ , and substituting in (D15) and then using (D16), leads us to the following proposition.

The general equilibrium productivity effect characterised in 3 implies that the employment effect is still subject to diminishing returns in general equilibrium. Therefore, robots are more likely to boost employment in a country at early stages of industrialisation.

E Calculation of Downstream Exposure to Robots

To construct downstream exposure to automation, we first assign the 9-digit codes to the corresponding 2-digit industry. Then, for each 5-digit buying industry, we compute the base-year share of expenditure on 2-digit selling industries. This procedure leads to the loss of observations as reliable input data is not available for all plants in all years. We obtain a detailed 5-by-2-digit input-output table based on disaggregated information. We use the table to compute downstream exposure to robots in each 5-digit industry. Let $\sigma_{l,i}$ be the share of inputs bought by (5-digit) industry l from (2-digit) industry i . We calculate downstream exposure to robots as:

$$ETR_{l,t}^d = \sum_i \sigma_{l,i} \left[\frac{R_{i,t}}{L_{i,t_0}} \times \overline{secondary}_{i,t_0} \right] \quad (\text{E1})$$

where $\overline{secondary}_{i,t_0}$ is the industry i average share of secondary education workers.

We add $ETR_{l,t}^d$ and the interaction $ETR_{l,t}^d \times ETR_{f,t}$ as additional regressors in $X_{f,t}$ of Equation (6). The coefficient of $ETR_{l,t}^d$ measures the impact of downstream exposure to robots on non-exposed plants, while the coefficient of $ETR_{l,t}^d \times ETR_{f,t}$ measures the impact of downstream exposure to robots on the exposed ones. The latter variable tests the extent to which these backward linkage effects are more intense for non-automating plants.

Given that the inter-industry linkages are computed at the 5-digit industry-level, the 2-digit industry-year fixed effects included in (6) are not sufficient to address the endogeneity concerns that would arise using Indonesian penetration. Therefore, we use a reduced form specification replacing $\frac{R_{i,t}}{L_{i,t_0}}$ with the OECD average penetration, which we interpret in terms of automation possibilities.

F Productivity Estimation

We assume that in each year t , plant f produces output Q_{ft} with the following production function:

$$Q_{ft} = \min \left\{ \gamma_m M_{ft}, F(K_{ft}, L_{ft}) \cdot \Omega_{ft} \right\} \quad (\text{F1})$$

where M_{ft} is raw material use, K_{ft} the capital stock and L_{ft} labor. The term Ω_{ft} represents Hicks-neutral productivity.

The production function (F1) is a structural value added specification [De Loecker and Scott \(2016\)](#) in which capital and labor are allowed to be characterised by some degree of substitution and materials use is a perfect complement to the combination of the other inputs.

Given (F1), a profit maximising plant sets

$$Q_{ft} = \gamma_m M_{ft} = F(K_{ft}, L_{ft}) \cdot \Omega_{ft} \quad (\text{F2})$$

To estimate production function parameters, we take the logged version of (F2):

$$q_{ft} = f(k_{ft}, l_{ft}; \beta) + \omega_{ft} + \epsilon_{ft} \quad (\text{F3})$$

Productivity ω_{ft} is known by plants' managers, but not by us. The variable ϵ_{ft} is an i.i.d. error term that captures disturbances such as measurement errors.

We are interested in estimating the vector of the production function parameters β . To recover unbiased and consistent estimates of firms' production function (F3), we need to address the well-known simultaneity problem deriving from the fact that ω_{ft} is correlated to labor but not to capital, which is chosen one period ahead. Thus, we build on the methodology of [Akerberg et al. \(2015\)](#) and make the following timing assumptions concerning inputs' decisions: i) capital k_{ft} is chosen at $t - 1$; ii) l_{ft} is chosen at $t - b$ after observing ω_{ft} , and iii) (log) materials use m_{ft} is chosen at $t - a$, with $1 < b < a$.

We then specify the materials demand function, $m_{ft} = \tilde{h}(\omega_{ft}, k_{ft}, l_{ft})$.

Assuming that the materials demand function is monotonically increasing and invertible in ω , we obtain a control function that proxies for unobserved productivity:

$$\omega_{ft} = h(m_{ft}, k_{ft}, l_{ft}) \quad (\text{F4})$$

where $h \equiv \tilde{h}^{-1}$. Adding $h(\cdot)$ to (F3), we get

$$q_{ft} = f(k_{ft}, l_{ft}; \beta) + h(m_{ft}, k_{ft}, l_{ft}) + \epsilon_{ft} \quad (\text{F5})$$

We follow [Akerberg et al. \(2015\)](#) in approximating the right-hand-side of (F5) with a third-order polynomial in all its elements. From the first stage, we obtain expected output \hat{q}_{ft} and the residuals $\hat{\epsilon}_{ft}$.⁵⁷

The next step is specifying a law of motion for productivity ω_{ft} . We assume that ω_{ft} follows a Markov process that can be shifted by plant managers' action:

$$\omega_{ft} = g(\omega_{f,t-1}, \Gamma_{f,t-1}) + \xi_{ft} \quad (\text{F6})$$

In (F6), ξ_{ft} denotes the innovation to productivity and the vector Γ includes variables controlled by plants' managers that influence the expected future value of productivity and state

⁵⁷ It should be noticed that in the first stage, none of the production function parameters are identified, because they enter both $f(\cdot)$ and $h(\cdot)$.

variables which determine differences in productivity dynamics across plants. In our framework, these variables capture the opportunities for automation available to each plant, which we proxy by plant-level exposure to robots as defined by (5).⁵⁸ Current expected productivity is then expressed as a function of the data and parameters:

$$\omega(\boldsymbol{\beta})_{ft} = \hat{q}_{ft} - f(k_{ft}, l_{ft}; \boldsymbol{\beta}) \quad (\text{F7})$$

To estimate $\boldsymbol{\beta}$, we form moments based on the innovation ξ_{ft} in the law of motion (F6):

$$\xi(\boldsymbol{\beta})_{ft} = \omega(\boldsymbol{\beta})_{ft} - E[\omega(\boldsymbol{\beta})_{ft} | \omega(\boldsymbol{\beta})_{f,t-1}, \mathbf{\Gamma}_{f,t-1}] \quad (\text{F8})$$

The moments that identify the parameters are:

$$E[\xi(\boldsymbol{\beta})_{ft} \mathbf{M}_{ft}] = 0 \quad (\text{F9})$$

where the vector \mathbf{M}_{ft} includes current capital, lagged labor, and lagged materials use.

In our empirical application, we use a flexible trans-log specification to approximate $f(\cdot)$. Our setup and timing assumptions are based on the idea that materials is the most flexible input in production. We deflate value added and materials expenditure with the plant-specific deflators (see Appendix B.6). For capital, we employ asset specific price indexes, distinguishing between machinery and equipment, vehicles, buildings, and land.

We obtain the production function parameter vector $\hat{\boldsymbol{\beta}}$ by estimating F9 with GMM and bootstrapping errors over hundred repetitions.

F.1 Deriving TFPQ and Real Marginal Cost From Plants' Cost Minimisation

Quantity-total factor productivity (TFPQ) is obtained using (F7).

We follow De Loecker and Warzynski (2012) to obtain a measure of plant level markup from the plants' first order conditions. Cost minimisation with respect to labor, which we

⁵⁸ In our application we impose a simple AR(1) form for (F6).

consider a static input, implies the following first order condition:

$$\frac{\partial \mathcal{L}_{jt}}{\partial L_{ft}} = W_{ft} - \lambda_{ft} \frac{\partial F(K_{ft}, L_{ft}) \cdot \Omega_{ft}}{\partial L_{ft}} = 0$$

where \mathcal{L} is plant's f Lagrangian, W_{ft} wages and λ_{ft} the Lagrangian multiplier. Rearranging terms and multiplying both sides of the previous equation by $\frac{L_{ft}}{Q_{ft}}$, we obtain

$$\frac{\partial F(K_{ft}, L_{ft}) \cdot \Omega_{ft}}{\partial L_{ft}} \frac{L_{ft}}{Q_{ft}} = \frac{1}{\lambda_{ft}} \frac{W_{ft} L_{ft}}{Q_{ft}}$$

As in [De Loecker and Warzynski \(2012\)](#), we define the plant's markup over the marginal cost of output λ_{ft} as

$$\mu_{ft} \equiv \frac{P_{ft}}{\lambda_{ft}}$$

where P_{ft} is the price of output produced by the plant. The previous equation yields an expression of plants' markup depending on the elasticity of output with respect to the variable input, β_l , and the inverse of the revenue share of expenditure on L_{lt} :

$$\mu_{ft} = \frac{\partial F(K_{ft}, L_{ft}) \cdot \Omega_{ft}}{\partial L_{ft}} \frac{L_{ft}}{Q_{ft}} \frac{P_{jt} Q_{jt}}{W_{jt} L_{jt}} = \beta_l \frac{P_{jt} Q_{jt}}{W_{jt} L_{jt}}$$

In our empirical application, the markup is given by

$$\mu_{ft} = \beta_l \frac{P_{ft} Q_{ft}}{W_{ft} L_{ft}} \frac{1}{\hat{\epsilon}_{ft}} \quad (\text{F10})$$

where the last term in (F11) is the residual obtained from the first stage estimation of (F5). As discussed in [De Loecker and Warzynski \(2012\)](#), including $\hat{\epsilon}_{ft}$ is important, as it allows to purge the estimated markup for variation in output not due to changes in inputs.

Finally, we recover marginal cost as

$$mc_{ft} = \frac{P_{ft}}{\mu_{ft}} \quad (\text{F11})$$